SINTEF

2019:01000- Unrestricted

Report

Simulation of multi-market trading

A case study

Author(s)
Marte Fodstad
Michal Kaut

SINTEF Energy Research
Production Planning
2019-09-25

SINTEF Energi AS
SINTEF Energy Research
Address:
Postboks 4761 Sluppen
NO-7465 Trondheim
NORWAY
Switchboard: +47 73597200
Telefax: +47 73590260
energy.research@sintef.no
www.sintef.no/energi
Enterprise /VAT No:
NO 939 350 675 MVA

Report

Simulation of multi-market trading

A case study

KEYWORDS:
Hydropower
Power markets
Bidding
Scheduling

VERSION
1.0

DATE
2019-09-25

AUTHOR(S)
Marte Fodstad
Michal Kaut

CLIENT(S)
MultiSharm project group

CLIENT'S REF.

PROJECT NO.
502000928

NUMBER OF PAGES/APPENDICES:
35 + Appendices

ABSTRACT
This report describes a study of day-ahead (DA) and intraday (ID) trade for a hydropower producer with a single reservoir and power plant. The study simulates the trade decisions and plant operation in a rolling-horizon setup where DA decisions are reoptimized on a daily basis and ID decisions are reoptimized on an hourly basis. Market prices are given by historic German prices from 2017, while price uncertainty is modelled by sampling from historical price paths using kernel-regression. The study assesses if the producer can gain from participating in the ID market in addition to DA trade. ID trading motivated by suboptimal DA commitments due to bid curve requirements or motivated by short-term production capacity uncertainty (e.g. inflow or outages) are not captured in the study.

Main findings are that the change in performance is relatively small when trading in both DA and ID market compared to only DA trade, and whether a gain or a loss is observed depends on the assessment interval. Even though ID trade is only conducted on a short horizon, within day, the end-of-horizon modelling (model horizon and water value description) can significantly affect the performance with ID trade.

PREPARED BY
Marte Fodstad

SIGNATURE
Marte Fodstad

CHECKED BY
Hans Ivar Skjelbred

SIGNATURE
Hans Ivar Skjelbred

APPROVED BY
Knut Samdal

SIGNATURE

REPORT NO.
2019:01000

ISBN
978-82-14-06372-1

CLASSIFICATION
Unrestricted

CLASSIFICATION THIS PAGE
Unrestricted

Document history

VERSION	DATE	VERSION DESCRIPTION
0.1	2019-08-28	Draft version –for review in project group. Not approved
0.2	2019-09-10	Updated draft based on comments from project group. Not approved
1.0	2019-09-25	First unrestricted version.

Preface

This report is written as part of the MultiSharm project, funded by Statkraft Energi, Hydro Energi, E-CO Energi, Agder Energi, TrønderEnergi and the Research Council of Norway (grant number 243964). The authors would like to thank the funders for supporting the project and for all fruitful comments on the way. Furthermore, we thank former colleague Kjetil T. Midthun and Trine Boomsma at the University of Copenhagen for active contributions in the design phase of this study. Lastly, several SINTEF colleagues have strengthened and facilitated the work through discussions, technical support and quality assurance, in particular Hans Ivar Skjelbred, Arild Helseth, Christian Øyn Naversen and Ellen K. Aasgård.

Table of contents

APPENDICES

[List appendices here]

The objective of this study is to compare the scheduling of a hydropower plant over a period of multiple weeks, preferably a year, trading in both day-ahead (DA) and intra-day (ID) relative to DA-trading only. A secondary objective is to build experiences on the model setup for such an analysis, in particular the dynamics when simulating several consecutive hours and days.

1 Market assumptions

The model represents a hydropower plant that is assumed to be price-taker in both markets. Since the plant do not have any pumping facility and no long-term contracted load obligations no purchase is possible in the DA market. In the ID market, on the other hand, both sale and purchase are allowed at the same price and net positions are reported. When modelling multiple markets in the same optimization model the model will seek to gain on any price difference between the two markets by taking opposite positions. We use sampling from historical prices in our scenario trees, and even though the markets on expectation over time have the same price, this might not hold for each of our individual scenario trees, which means arbitrage opportunities will exist. To limit the utilization of arbitrate possibilities the ID position is limited by the production capacity. An alternative approach would have been to assume a certain price elasticity, but this would contradict the price-taker assumption.

The ID market is modelled as if it was a cleared market with a single hourly price and hourly trade. This reduces the complexity of the model compared to modelling the real continuous ID trading which would require modelling of the price and quantity development of ID bids when approaching ID gate closure. Our rough assessments of this simplification indicates that the value achieved in the single-price approximation does not exceed the value that can be achieved in the continuously traded market, since the continuous version gives the possibility to accept or reject bids on multiple price levels, giving stronger possibilities for price differentiation than what can be achieved with a single price. These assessments did not take any quantity limitations on the market side into account.

2 Case setup and information structure

The simulation is run for three market combinations as illustrated in Figure 1:

- DAonly – optimizing DA sale without any ID market knowledge
- DA-ID – taking the DA commitments from DAonly as a given load, optimizing ID sale
- DA+ID – integrated optimizing of both DA and ID sale.

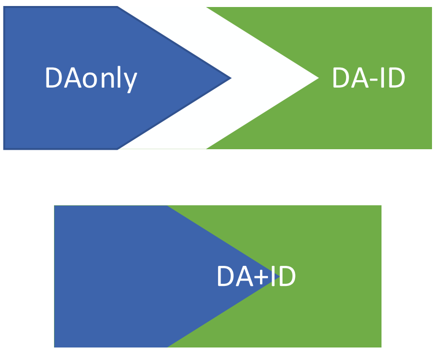

Figure 1 The three simulated market combinations. DAonly, ID with a given DA posisiton (DA-ID) and DA and ID integrated in one model (DA+ID)

This makes it possible to assess

1. the added value of multi-market participation (DA-ID or DA+ID vs DAonly)
2. the added value of integrated/coordinated trading vs sequential (DA+ID vs DA-ID)

It should be noted that the setup does not distinguish between DA bidding and DA scheduling after the market clearing.

The simulation setup seeks to replicate the market trading steps of daily DA commitments, at noon, and hourly ID commitments by running the SHARM model ones for every trading step, as illustrated in Figure 2. Each of these model runs are denoted subproblems. A subproblem optimizes and records the operation of an hour or a day, but the model horizon is longer than the recorded period to reduce end-of-horizon effects. Unless anything other stated, references to the results of a subproblem or overall simulation means the results from this recorded period and not the whole model horizon.

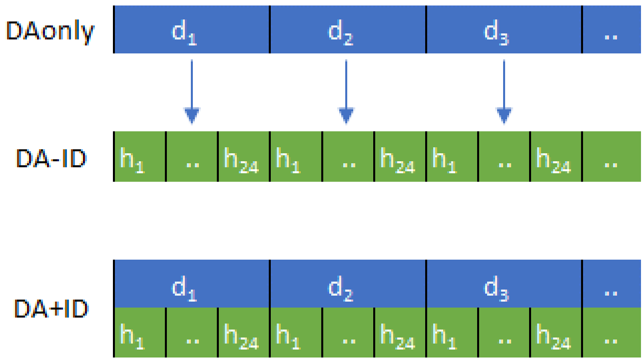

Figure 2 Step length between each subproblem in the simulation of each market combination.

Within the simulation of one market combination all subproblems are run in a consecutive sequence. The end reservoir level and generator state of one subproblem are the initial state of the next. A new scenario tree is used for each subproblem, while water values are updates once a week.

The real-world multi-market structure gives a vast number of decision stages, at least 25 per day for DA and each of the 24 hours of ID trade. To be able to solve the subproblems for each time step through the year within reasonable computation time, some trade-offs between modelling precision and problem size are needed. In each subproblem the price of the currently traded hour(s) is deterministic and equal to the historically observed price, while future prices are stochastic and based on sample-paths of prices from history, as described in Section 3.2. The scenario trees are 2-3 stages, and the modelling horizon is kept short relative to the common practice with SHOP. The time and stage structure of the subproblems is different for the different market combinations DAonly, DA-ID and DA+ID and will be described in the following. When setting up the subproblem structures it has been a goal to support a fair comparison between coordinated and sequential trading.

2.1 Subproblem structures

DAonly has a 3-days rolling horizon, with a one-day deterministic period (first stage), as illustrated in Figure 3. The DA price is assumed deterministic the first day by assuming that the producer would be able to perfectly fit the bid curve to the marginal cost curve, and thereby match the quantity to the cleared market price. This setup is equivalent to the real market setup where day-ahead decisions are taken at 12 noon, since the DA subproblem do not cover any uncertainties or decision taken between midday and midnight.

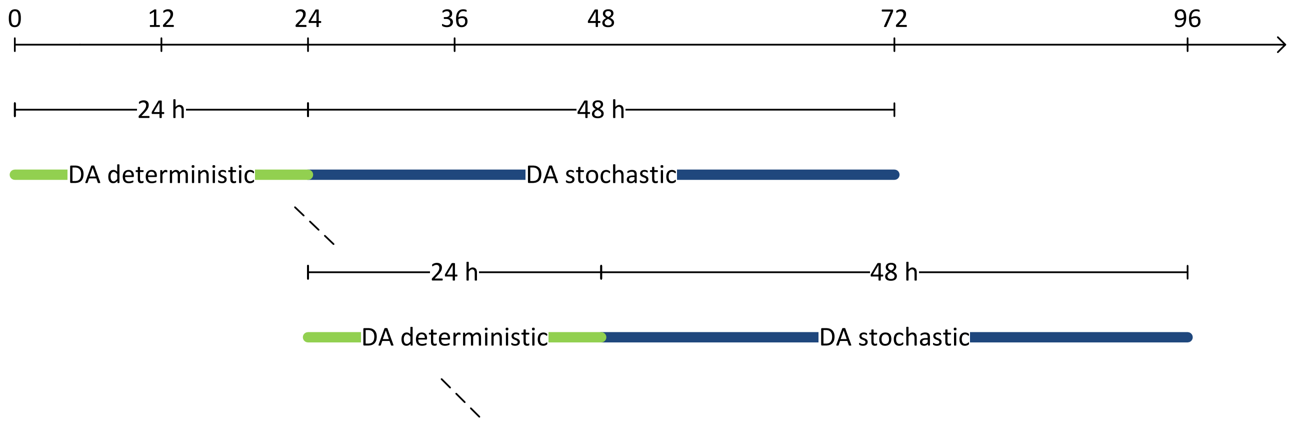

Figure 3 DAonly daily subproblem

DA-ID optimizes ID trade and operation for the DA load given by preceding DAonly run. The subproblem is run with hourly steps over a 12-24 hour horizon. When optimizing hour 0 (midnight until 1 am) the model horizon is 24 hours (the whole day). For each hourly subproblem the horizon is one hour shorter until hour 12, when next day's DA commitments becomes known and the model horizon is extended with 24 hours. The first hour is deterministic in each hourly subproblem.

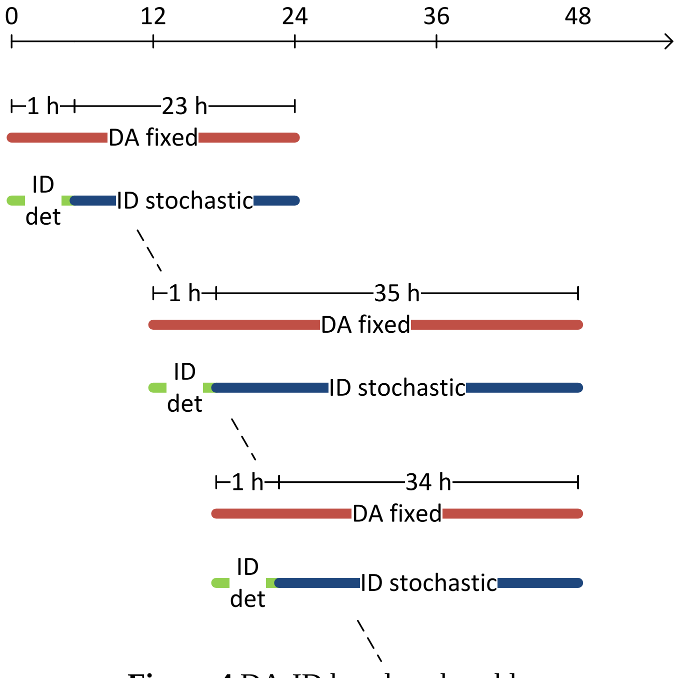

Figure 4 DA-ID hourly subproblem

DA+ID optimizes DA and ID decisions jointly with an hourly updated optimization to simulate the ID market in a similar way as for sequential optimization. The model horizons will vary over the hours of a day, as illustrated in Figure 5. ID will have a 13-36 hour horizon, always ending the horizon at midnight, with one new day added at 12 noon. ID has deterministic price the first hour and stochastic price the rest of the hours. DA is fixed by earlier model runs for the current day, and for any hour after 12 noon DA is also fixed for the coming day. In the optimization representing 12 noon DA is deterministic in the coming day, and the end model horizon is rolled forward to cover one more day. For all other problems than 12 noon there are no DA decisions, and the fixed DA horizon corresponds to the ID horizon.

The scenario tree is a two-stage tree with branching caused by ID prices in all hours except 12 noon. At 12 noon the tree is three-stage with a one hour first stage, ID prices causes the branching into the second stage and DA prices causes the branching into the third stage.

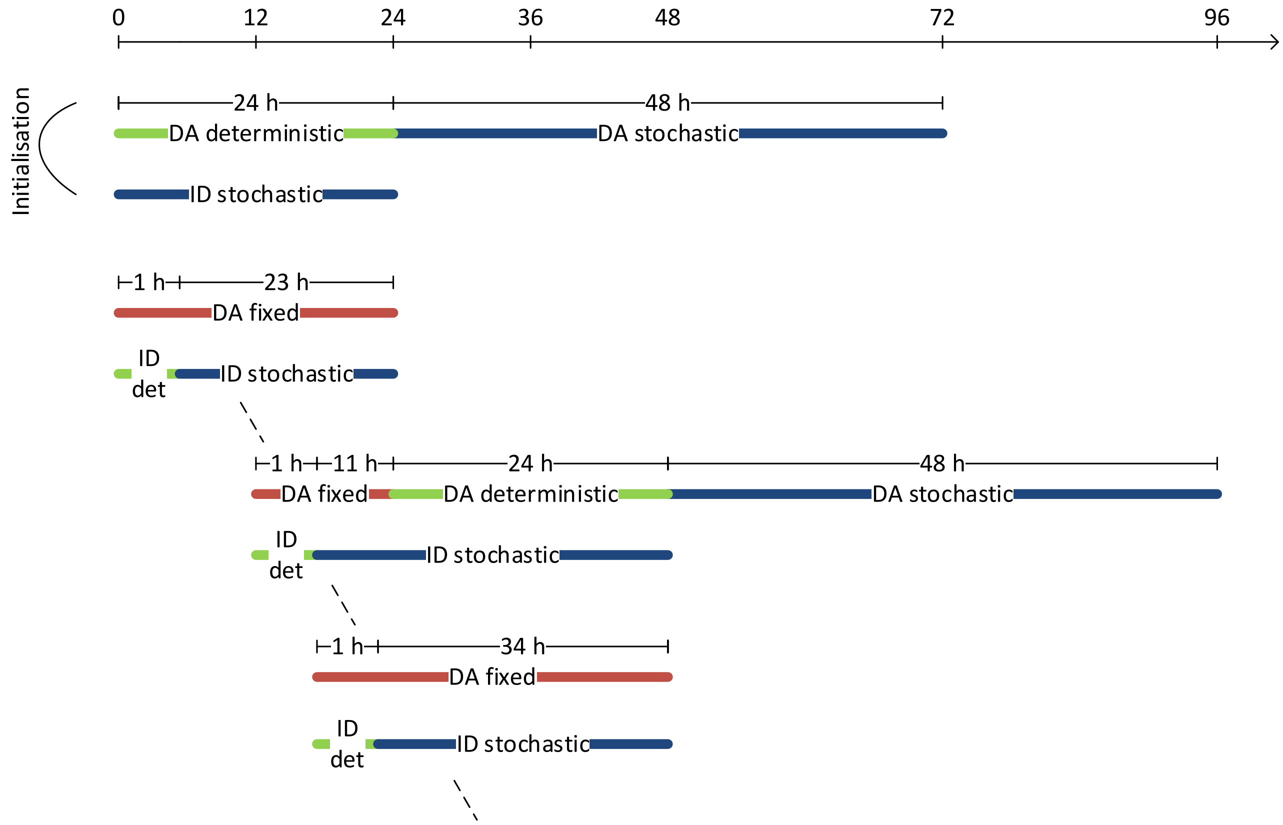

Figure 5 DA+ID hourly subproblems illustrated for three different hours of the day

2.2 Subproblem structure with extended model horizon

Motivated by simulation results presented in Subsection **Error! Reference source not found.** an alternative subproblem structure is also tested. In this setup the model horizon of those subproblems not making DA decisions are extended relative to what is described in Subsection 2.1. This applies for all DA-ID subproblems and DA+ID subproblems not starting at noon (h==12). The extension is a 48 hours increase of the model horizon. This makes 61 hours the shortest model horizon and the DA-ID and DA+ID subproblems end at the same hour as the corresponding DAOnly subproblems. Within this horizon extension only DA trading is allowed. The DA price is uncertain in this period, creating a new decision stage in the scenario trees. This turns all scenario trees into three-stage trees, similar to the trees originally used for DA+ID at noon.

In the following presentation of the scenario generation procedure the original subproblem structure is assumed. The same procedure is used to generate scenario trees also for the extended subproblems. When presenting simulation results the original subproblem structure is assumed unless otherwise is stated.

3 Price modelling

In the following subsections, we describe the procedure for generating power-price scenarios, plus results of data analysis required to justify the procedure.

3.1 Data analysis

The data analysis is based on German day-ahead and intra-day prices from period 2013–2017. We present only the main findings, the full analysis is presented in a separate document "*Analysis of German day-ahead and intra-day electricity prices*". The day-ahead prices are presented in Figure 6.

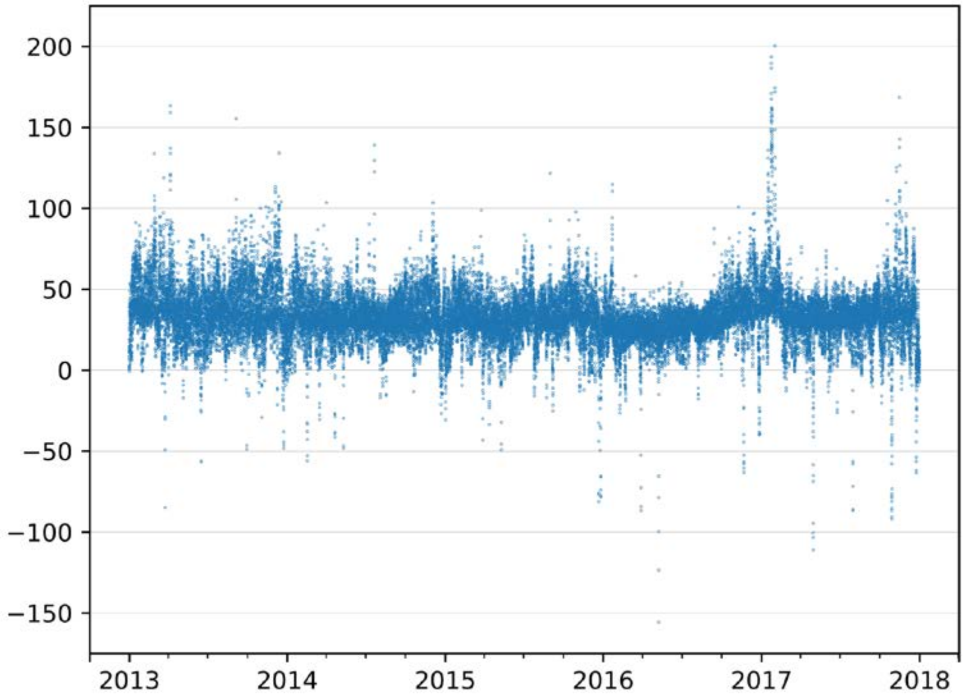

Figure 6 Historical day-ahead electricity prices, in €/MWh

The time structure of the model implies that we need to investigate two sets of distributions:

- distribution of intra-day prices, given day-ahead prices for the same day,
- distribution of day-ahead prices, given day-ahead and intra-day prices for the previous days.

We start with the intra-day prices, conditional on day-ahead prices for the same day. To handle the dependency, we study the distribution of the difference ID – DA, which we will refer to as the *intra-day price correction*. First, we check whether these corrections depend on the day-ahead prices: if they do, we would need to handle the dependence in some way. On the other hand, should the price corrections be independent on day-ahead prices, it would simplify both the data analysis and the scenario generation as we would need to estimate and model only one, global, distribution for the price corrections.

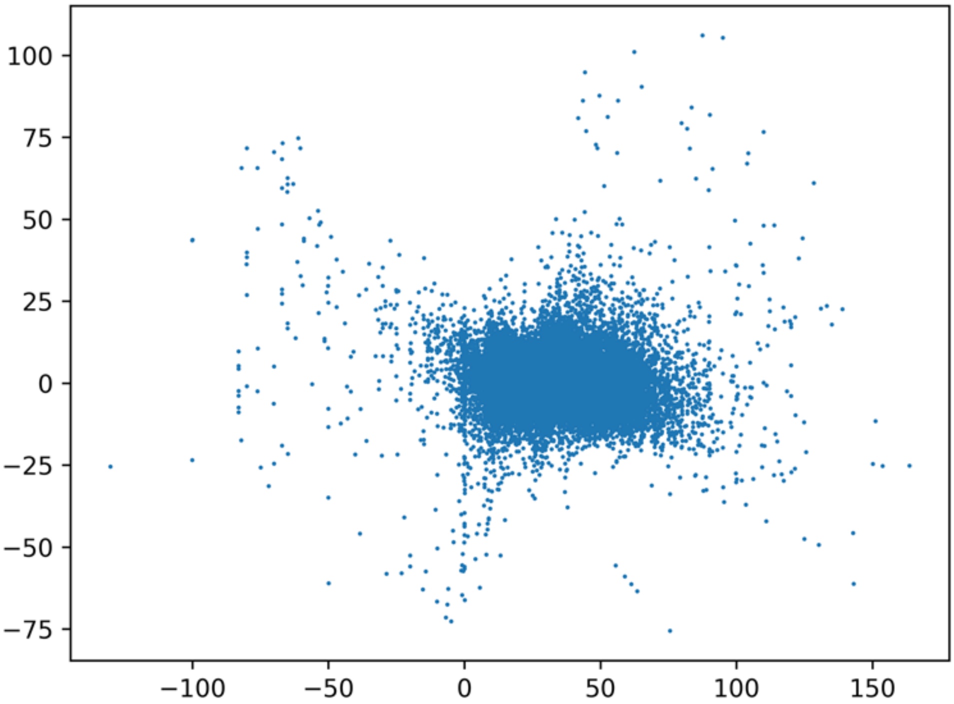

Figure 7 Intra-day price corrections (vertical axis) vs day-ahead prices (horizontal axis)

Fortunately for us, the prices do not show any sign of dependence, as can be seen from the scatter plot in Figure 7 and also supported by a negligible correlation of -0.07. The same holds if we split the data and only consider one season at a time or one week-day at a time, as documented in the aforementioned memo.

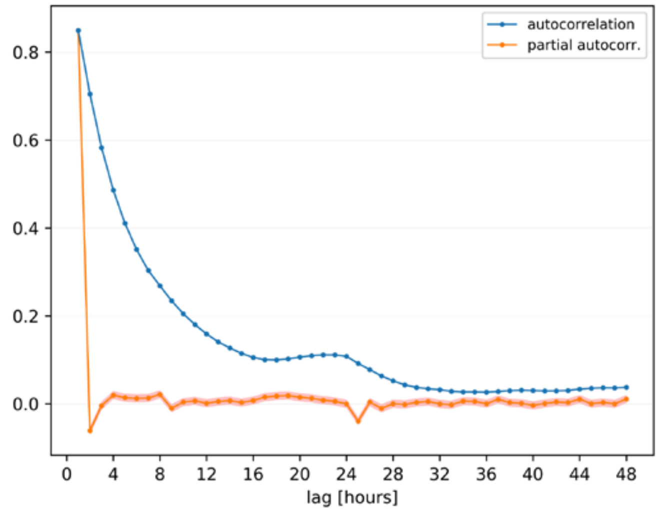

Figure 8 Autocorrelations of the intra-day price corrections

Now we know that to generate scenarios for the intra-day prices, we need only to replicate the distribution of the intra-day price corrections, given price corrections in the previous hours. The next question is how many hours back do we need to look, which is connected to autocorrelations of the price-correction series. This is shown in Figure 8, where we can see that the only substantial partial autocorrelation is for lag equal to one, implying that the distribution of price corrections starting from hour *h* is mostly determined by the price correction at hour *h*-1.

The next step is to test the distribution of day-ahead prices. This is required for the scenarios for the 12th hour, where we build a scenario tree with uncertain intra-day prices for the next day (hours 13-36 in the model) and uncertain day-ahead prices starting the day after that, i.e., after hour 36. This implies that we need to know the distribution of the day-ahead prices, conditional on both the day-ahead and intra-day prices up to the starting day.

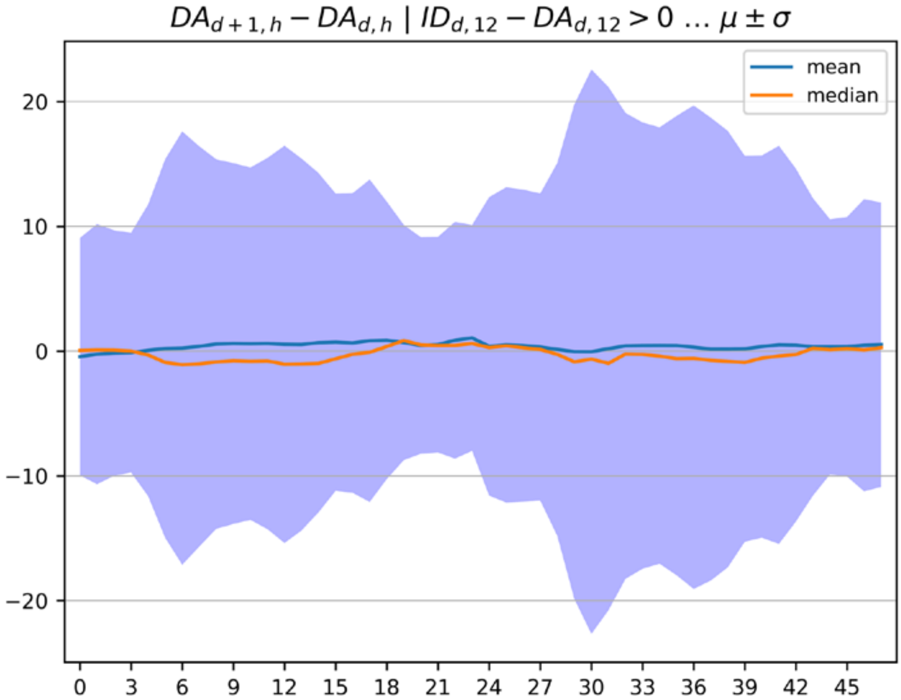

Figure 9 Distribution of the daily changes in the day-ahead prices, conditional on positive intra-day price reductions. The blue band is mean ± one standard deviation

Just as in the previous case, it would simplify the scenario generation if the day-ahead prices were independent on the intra-day prices from the previous day, conditional on the day-ahead prices for the same day—or, equivalently, independent on the previous-day's price corrections.
Note that it would be natural to expect some dependence here: if we had positive intra-day price correction yesterday, it means that prices yesterday were higher than we had thought, which could imply that prices for today would go up. To test this hypothesis, we checked the distributions of *inter-day changes* of the day-ahead prices, conditional on the price-correction in hour 12 of the previous day being positive; we use the 12th hour because that is the time where today's prices were decided. As we can see in Figure 9, both the mean and median of the resulting distribution is very close to zero, suggesting that there is no such effect.
In the memo, we have in addition compared several regression models for day-ahead prices, with and without the intra-day prices, and found that including the intra-day prices does not improve the predictive power of the models. This is also an evidence for independence.

3.2 Scenario generation

The scenario-generation procedure is based on a kernel-regression approach from (Pflug & Pichler, 2016). They show how to estimate a conditional distribution from a set of paths, by weighting the paths by their similarity and "smoothing" the weighted observations using a kernel.
While we could use the approach directly, it would mean generating new, synthetic, observations. In that case, it would be difficult to guarantee that the dependencies between hours are realistic. Instead, we stop

with the weights and then simply use the observations from the S days with the highest weights, with probabilities given by the weights scaled to sum up to one. In the context of the original method, this roughly corresponds to using degenerate (Dirac delta) kernel.

First, we start with the intra-day prices. In hour h, we are generating scenarios for intra-day prices for hours $h+1,..,T$, where T is 24 if h is smaller than 12, and 48 otherwise. Since our data analysis showed that we can generate scenarios for the intra-day price *corrections*, and that these are mostly only dependent on the current value of the price correction, the process to generate S scenarios is as follows:

1. Calculate weights for all eligible days in the dataset, based on the difference in their intra-day price correction at hour h. In our case, we define eligible days as the same day of week and same part of year, that is at most 60 days from the current day.
2. Select the S days with the highest weights
3. For each day, compute the intra-day prices as $ID_{h'} = DA_{h'} + PC_{h'}$, for all hours $h' > h$
4. These values form the scenarios, with probabilities equal to the weights, scaled so that they sum up to one

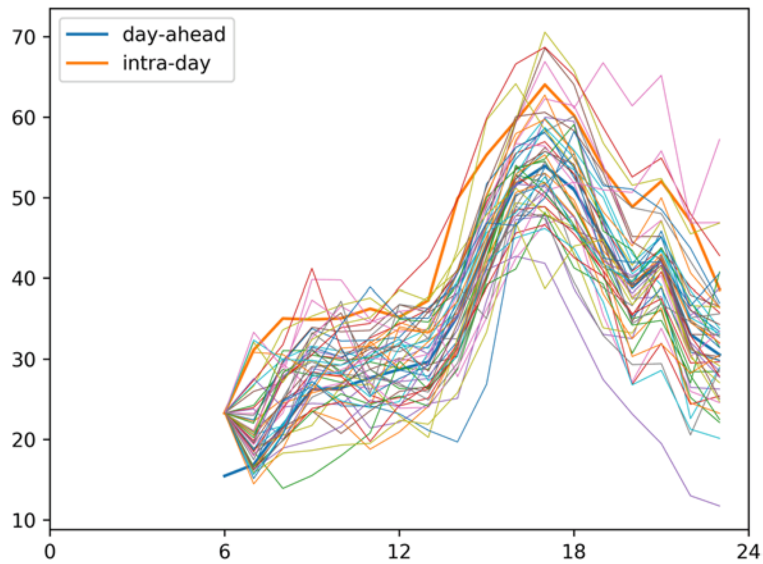

Figure 10 Example intra-day price scenarios with $h = 6$

An example of scenarios generated using this approach, using the German data, is in Figure 10.

Next, we move to scenarios for the day-ahead prices. These are generated at $h = 12$, but with stochastic DAs generated at $h = 48$ for two days ahead, i.e., with stochastic DAs at hours 49 to 96. The data analysis has shown that the distribution of DAs does not depend on the intra-day price corrections, but this still leaves us with dependency on the DAs at previous hours.

We concentrate on the last 24 hours, i.e., want to find days in the historical data with DAs similar to what we have in hours 25 to 48. Since the prices are strongly autocorrelated, we do not need to compare all 24 hours. Hence, we base the comparison on values at hours 25, 31, 37, 43, and 48; that is, values 1, 6, 12, 18, and 24 hours before the starting $h = 49$. The scenario-generation process is the same as for the intra-day prices, so we do not repeat it here.

An example of a complete scenario tree for $h = 12$, using the German data and including scenarios for both DAs and IDs, is in Figure 11.

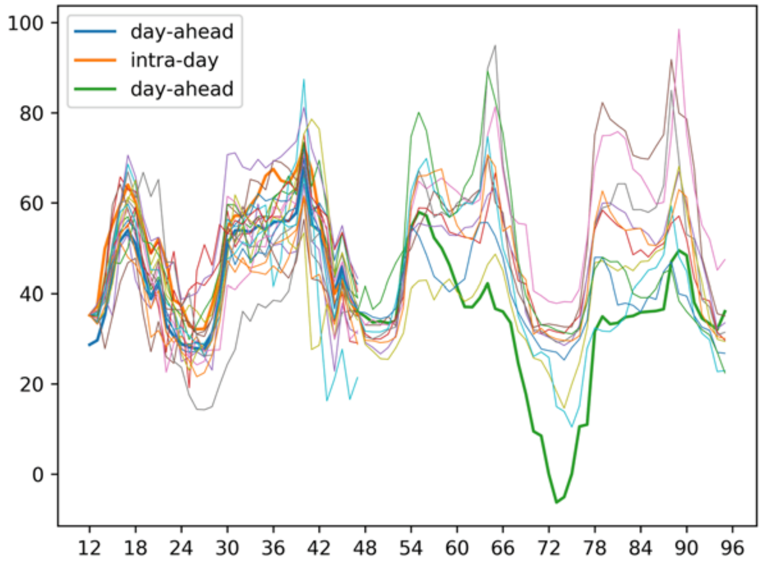

Figure 11 Scenarios hour 12, including scenarios for both day-ahead and intra-day prices

The results presented later in this document are based on scenario trees with 10 samples for ID and 6 samples for DA, giving 60 scenarios in the trees combining DA and ID.

4 Production system modelling

Data representing a real production system with a single reservoir and two generators is used in the study. Main characteristics of the production system is given in Table 1. As can be observed, the reservoir is small relative to the installed capacity and inflow, which indicates a relatively inflexible hydropower plant with mainly short-term balancing capabilities.

Table 1 Main characteristics of production system

Property	Value
Initial reservoir volume	82%
Difference between highest and lowest regulated level	2.7 m
Total max power output	Ca 90 MW
Full load hours[1]	42%
Time to empty a full reservoir[2]	60 h
Avg. time to fill an empty reservoir	144 h

Hourly inflow data for 2017 from the Norwegian regulator, NVE, are used. This inflow series represents one of the major catchment areas supplying the modelled reservoir and is scaled to match the average yearly inflow to the reservoir. Water value calculations in ProdRisk use openly available data from NVE included in SINTEF's inhouse dataset for the modelled watercourses. The data covers 50 years of daily historical inflow. The inflow data is plotted in Figure 12.

ProdRisk is run with inflow uncertainty, while we assume deterministic inflow in SHARM. In the long run inflow uncertainty is assumed to be substantial and important for the water value calculations, while the inflow uncertainty is assumed to be less important on a 36-12 hours horizon (from day-ahead bidding until

[1] Full load hours = Yearly inflow / (Max discharge * Hours in the year)
[2] Assuming no inflow

operation). Including inflow uncertainty will require data to describe this uncertainty and increase the model size and thereby the computational burden.

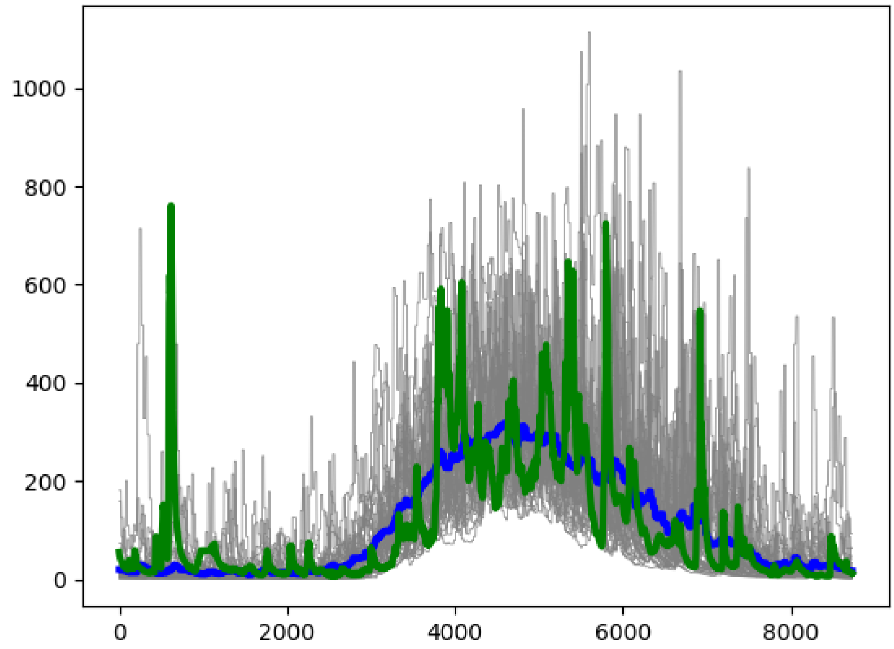

Figure 12 Hourly inflow for 2017 (green), daily inflow for 1958-2016 (gray) and average daily inflow (blue), in m^3/s

4.1 Water value calculations with ProdRisk

The history of price data is short, in particular for intra-day markets. To avoid limiting the history of inflows by the price data availability, we assume independence between the prices and inflows, and model a single average price in ProdRisk. This is clearly a simplification, but by SINTEF's experience such a single price usually gives more robust and consistent water value results than generating a Markov-model for prices in ProdRisk based on few price scenarios.

ProdRisk can only have one energy price pr hours, and we use the DA price. Data analysis of the German market data indicates that the price premium of ID relative to DA is on expectation zero, which makes the use of a single price less troublesome in a model with weekly granularity on the uncertainty. We run ProdRisk with an hourly time resolution to capture the price variations. This is important to achieve reasonable water values. Initial tests with daily resolution indicated too low water values with SHOP generally running on low reservoir levels.

The model horizon of ProdRisk need to be longer than for SHARM, typically 3 years, to be insensitive to the choice of water values describing the end-of-horizon in ProdRisk. Due to the relative small reservoir in the test case, less than three years would have been sufficient, but for simplicity the whole 3-year horizon is used. These ProdRisk end-of-horizon water values are established manually as a function with shape established for an other case and scaled to fit this study's price level. The extended model horizon in ProdRisk implies that price data for ProdRisk is needed after the year of study, 2017. Due to the lack of availability of a good quality price forecasts after 2017, we replicate the 2017 prices for another two years as price input to ProdRisk. This gives a somewhat abrupt price change at new year where the price series is concatenated, but this change is not larger than price changes found within the price series, and are therefore found acceptable in a rough approximation.

ProdRisk provides water value functions for the end of each week due to its weekly resolution on uncertainty. We assume that these water value functions are valid for the whole week, so that the water value function calculated for Sunday are used for all days of that week.

A possible improvement of the study would be to increase the resolution of the water values to daily. Three alternative approaches were assessed but dismissed due to work and/or computation load: The first alternative would be to assume a linear interpolation between the water value functions from one Sunday to the next. This will give a smoother transaction between days, but on the other hand requires some efforts to calculate correct linear interpolations for the multi-dimensional water value functions. The second alternative is to run ProdRisk repeatedly rolling forward one day for each model run. In addition to the added computational time this requires rewriting the iteration logic and data handling in the simulator. The third alternative is to extend the SHARM model horizon to reach the next Sunday, which naturally increases the computation time in SHARM.

5 SHARM run setup

SHARM is solved with four iterations for each subproblem, two full-iterations and two incremental iterations. The last full iteration included binary variables. SHARM parameter settings is given in Table 2.

Table 2 SHARM parameter settings

Parameter name	Parameter value
power_head_optimization	on
simple_pq_recovery	on
plant_uploading	lp
power_loss	pq
method	baropt
mipgap	0.005
timelimit	86400 for MIP-iteration, default (900) for other iterations

6 Results

In the presented results sales quantities and income are represented with positive values and purchase with negative values.

6.1 Model horizon

The results in Table 3 show the sum of accumulated net profits from simulations for 2017 (364 days) corrected for differences in end reservoir value for each of the market combinations.

Table 3 Achieved profit over the simulated year

Market combination	Total profit relative to DAonly
DAonly	100%
DA-ID	99.2%
DA+ID	98.8%

Somewhat surprisingly, the possibility to trade ID causes a loss in profit relative to DAonly both with sequential and integrated trading. It is tempting to think that ID trading is an option that *can* be utilized to improve the DA position but is not mandatory and would be left unused unless profitable. This is in theory a valid argument, *but*, to assess profitability in this context a perfect foresight on the consequence of ID trade is required. This means a perfect foresight assessment of the value of changing the reservoir levels and generator states when making ID decisions. In our rolling horizon simulation this is clearly not the situation, as the model horizon for ID trade in most subproblems is limited to the rest of the current day.

In the following sections we first show in detail a situation causing a substantial loss with ID trade compared to DAonly. Next, a discussion on possible measures related to end valuation of system states is discussed.

6.1.1 Example of loss in ID

This subsection describes an example of a situation where ID trade conducted after the DA trade (sequential) causes a loss compared to DAonly. This example is chosen because of its strong effect on the overall results, and not because it is a particular common situation. While the example is explained for DA-ID only, similar observations are found for DA+ID.

A plot structure, as in Figure 13, is used when presenting how the production system is operated under the different market assumptions. Inflow, discharge, bypass and spillage is plotted according to the left axis [m^3/s], storage level follows the first right axis [Mm^3], and the realized price is at the second right axis [EUR/MWh]. QMax, on the left axis, refers to the maximal discharge capacity of the plant. Note that the axes can have different scale in different plots.

As can be observed in Figure 12, there are some large inflow spikes, the first and largest in week 4 in 2017. This spike has inflow above the discharge capacity of the power plant, and over the duration of the spike the inflow even exceeds the reservoir capacity so that bypass or spillage is unavoidable. This is easily seen in Figure 13 representing the DAonly operation, where the reservoir is emptied in front of the spike[3] and discharge is at its maximum during the spike, but still some water is bypassed.

The next figure, Figure 14, shows that the DA-ID has a substantially larger bypass during the spike because the reservoir is not emptied prior to the spike. Clearly this means lost sales and a lost income potential. The difference in reservoir profile between DAonly and DA-ID mainly arises in the preceding week, shown in Figure 15 and Figure 16. DAonly plans substantial production through the week, but ID prices are below DA prices and also below the water values motivating DA-ID to buy back parts of the DA commitments. In this period none of the subproblems have sufficient model horizon to observe the spike.

The break point in the reservoir profiles around hour 528 indicates the first time the model horizon is long enough to include parts of the inflow spike. In the period from hour 528 to hour 576 the position of DA-ID worsens further relative to DAonly because the model horizon of DA-ID is shorter than that of DAonly, so that DA-ID still does not observe the full inflow spike and continues to buy back DA-commitments when ID-prices are low.

[3] Remember that the short-term subproblems in SHARM do not represent inflow uncertainty which makes the model more willing to extreme operation of the reservoir.

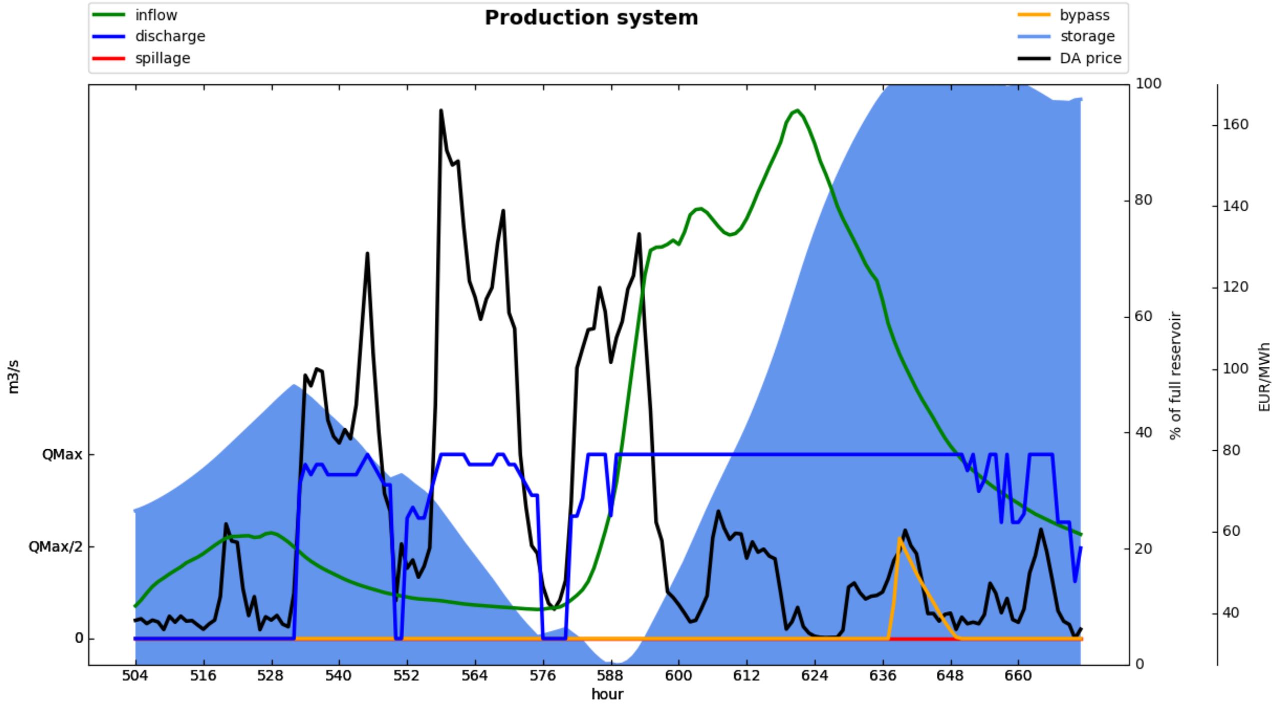

Figure 13 DAonly in week 4

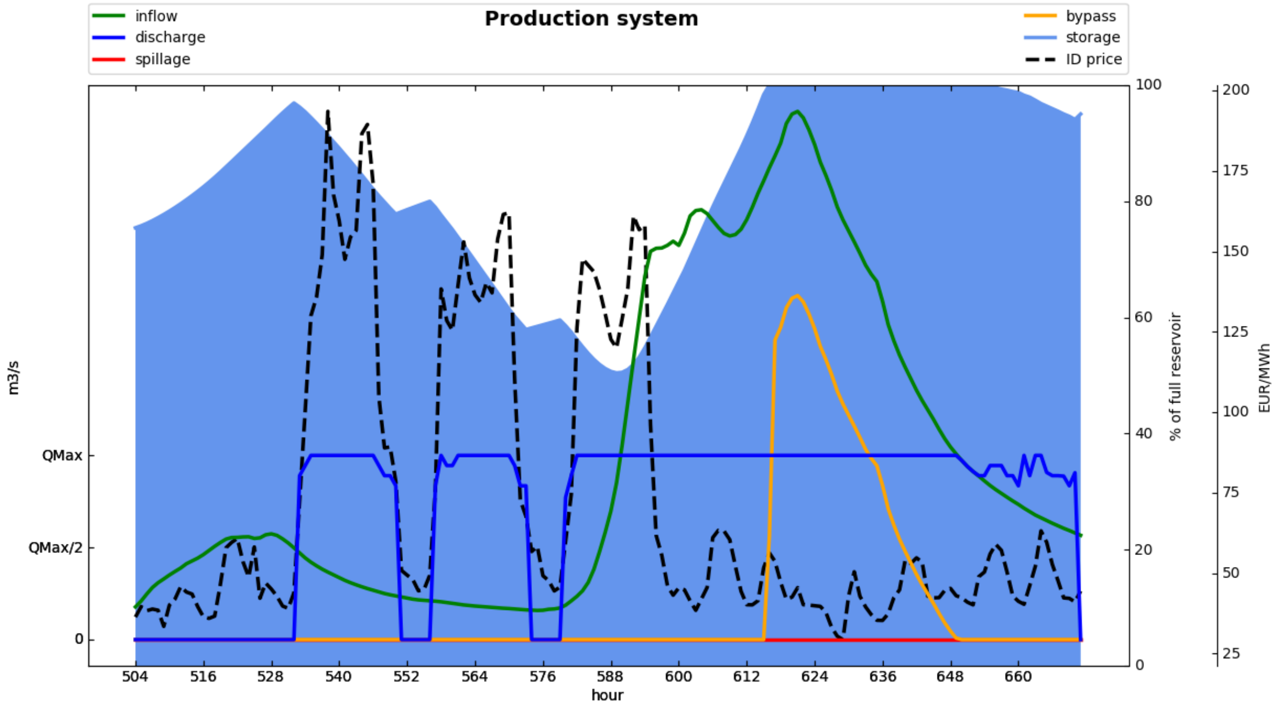

Figure 14 DA-ID in week 4

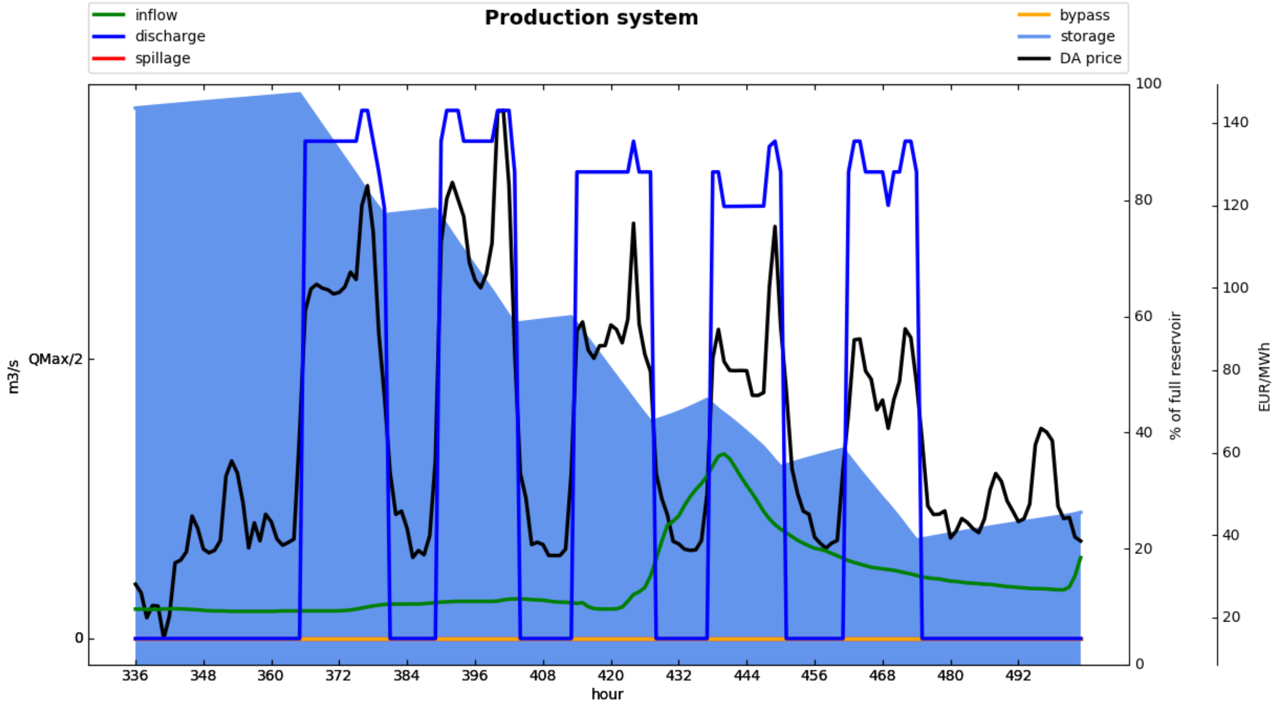

Figure 15 DAonly in week 3

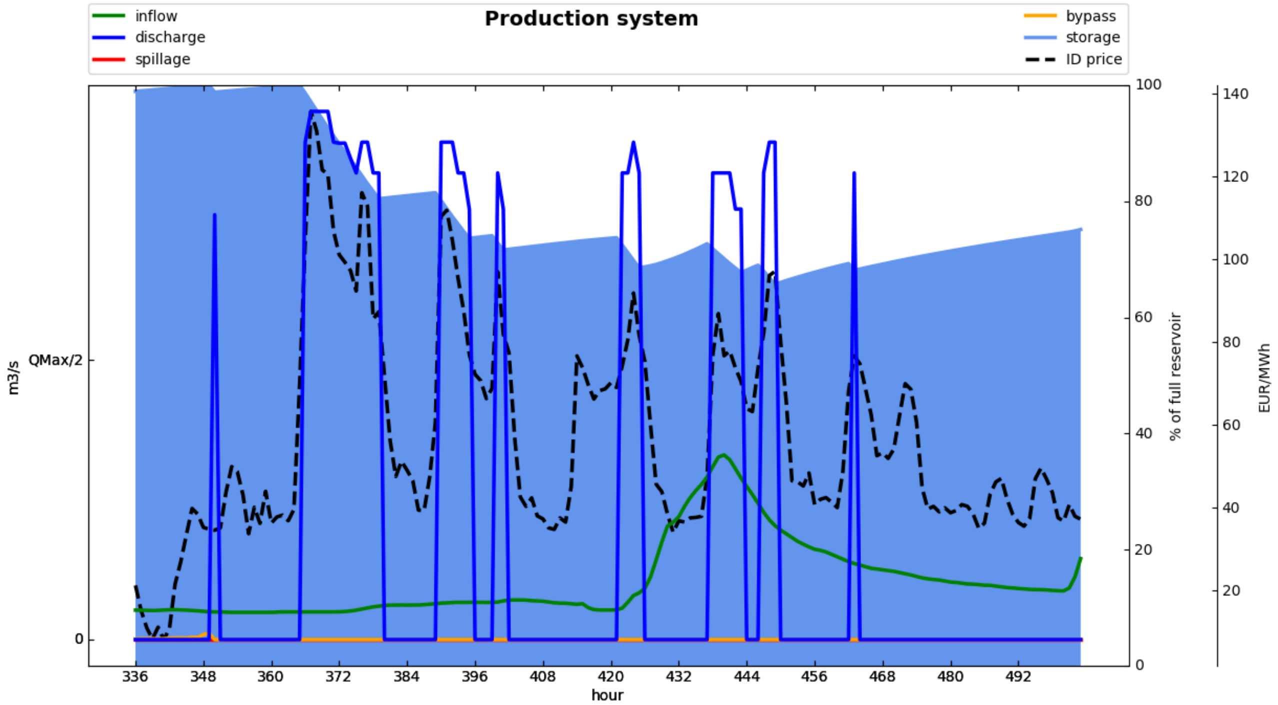

Figure 16 DA-ID in week 3

6.1.2 End valuation

The example in the previous subsection indicate that an improved end valuation might improve the performance of ID-trading. This is tested empirically in this subsection, utilizing a shortened test period representing week 1-4 in 2017.

To exclude other possible sources of differences between DAonly and DA-ID the following changes was introduced in the benchmark setup for this part of the study:

1) Only one full iteration with MIP was run for each subproblem avoiding the heuristic part of SHOP, but also removing the hear optimization functionality
2) Removed the uncertainty within the subproblem horizons. That is, the scenario trees with sampled prices was replaced with observed prices without any uncertainty. Still, the limited horizon naturally limits the perfect foresight.
3) The period of DA commitments (load) in DA-ID was extended to cover the whole subproblem horizon.

Over the 4-week period DA-ID reached a grand total at 87.5% of the DAonly grand total for the benchmark setup.

As DAonly and DA-ID use the same water value cuts from ProdRisk the shorter model horizon in DA-ID is assumed to be a drawback. The blue plots in Figure 17 show the DA-ID results as the model horizon is increased. The horizontal axis denotes the model horizon measured as the shortest horizon among the 24 subproblems within a day, this is the model horizon at 11:00. The shortest horizon, 13h, corresponds to modelling current day only, as described in Section 2. 61h corresponds to the same end hour as DAonly, while 73h means DA-ID never has shorter model horizon than DAonly. The results show that DA-ID need a longer model horizon to improve the grand total beyond what is achieved in DAonly. This happens at 88h minimum model horizon. With this model horizon length the end hour ends in the next week for some decisive hours in week 3 giving a stronger incentive to reduce the reservoir level, as shown in Figure 18, and thereby reduce bypass in week 4. The DA-ID gain is counterintuitively reduced when extending the horizon from 91h to 97h. This reduction, at 0.08%, is less than the MIP-gap tolerance when looking at absolute numbers for end value rather than end value relative to DAonly and is therefore assumed not significant.

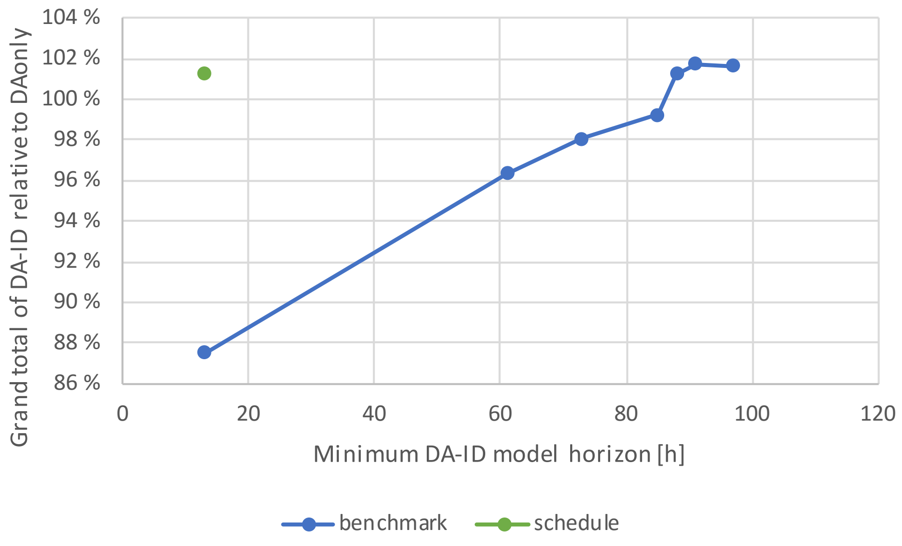

Figure 17 Sensitivity analysis of DA-ID performance

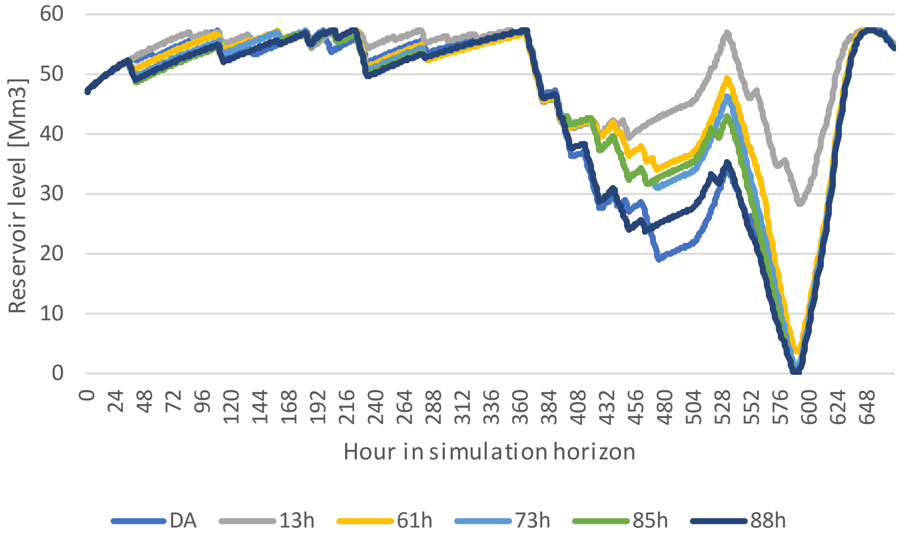

Figure 18 Reservoir paths

The motivation for initially giving DA-ID a shorter model horizon than DAonly was that DA is seen as the main market while ID was assumed used for repositioning within day. Following this idea, using the optimized reservoir path from DAonly to guide the DA-ID subproblem instead of using water value cuts is an option. This is implemented in SHARM by setting reservoir schedule on the end reservoir level.[4] The green plot denoted "schedule" shows the result with this approach, showing a gain from ID-trade within the current-day model horizon. A drawback of this approach is that computation time is more than 11 times the computation time of the benchmark with 13h horizon and more than twice the time for benchmark with 88h horizon.

6.2 DAonly, DA-ID and DA+ID profitability

This section presents results on the profitability of the different market combinations, both for the original subproblem structure and when subproblems have an extended model horizon as described in Subsection 2.2. Results with extended model horizon, identified by the extension "ext", is only available for parts of the model horizon[5] due to time consuming simulations and limits in time availability. To compare profitability of the different market combinations we use "profit", defined as total market income less costs. Cumulative profits are the profits accumulated from the start of the year.

From Figure 19 we observe that the impact of the different trading modes is relatively small compared to the main income trend driven by the inflow seasonal profile. To observe differences between the trading modes we use DAonly as a benchmark and plot relative values in Figure 20. Several values in early hours are outside the plotted range of 50%-120%. This is because we plot values relative to the DAonly cumulative profit which is small in the early hours. Therefore, even small deviations from the trading in DAonly will give large effect in the plot in early hours.

[4] SHARM do not support setting reservoir schedule at the end of the last mode hour, so the schedule is implemented at the second last hour.
[5] 6073 hours for DA-ID ext and 3086 hours for DA+ID ext

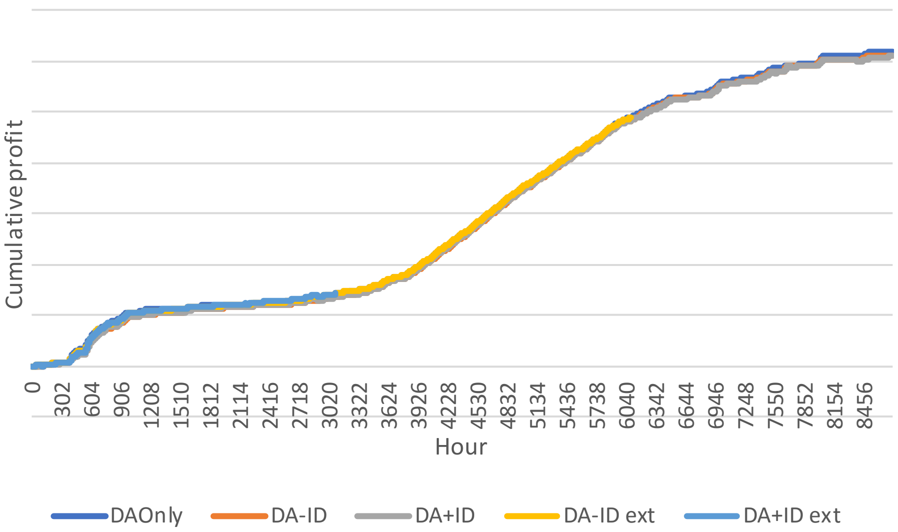

Figure 19 Cumulative profits

Furthermore, in the period until hour 672 (the first four weeks) the cumulative profits are dominated by the effect discussed in Subsection 6.1.1 where the simulations including ID withholds more water in the reservoir and therefore earns less in the market. Zooming in on this time period, in Figure 21 and Figure 22, we observe that the extended model horizon has the expected effect, that is, more income and lower reservoir level than the original model setup. The importance of this modelling choice is reflected in the long-lasting difference between the original and extended model in Figure 20. In Figure 21 we can also observe that DA+ID has a lower cumulative profit than DA-ID for both model horizons. This is a result of the increased freedom of DA+ID compared to DA-ID, as the former can adapt the DA sale to an expected ID position. Knowing, from Subsection 6.1.1, that the ID prices are low we see how the superior freedom are used to produce less.

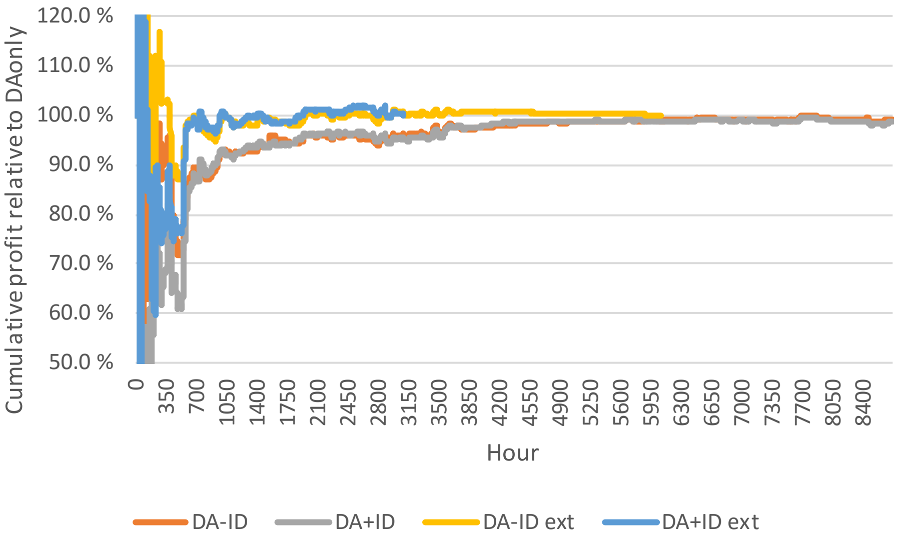

Figure 20 Cumulative profits with ID-trade relative to DAonly

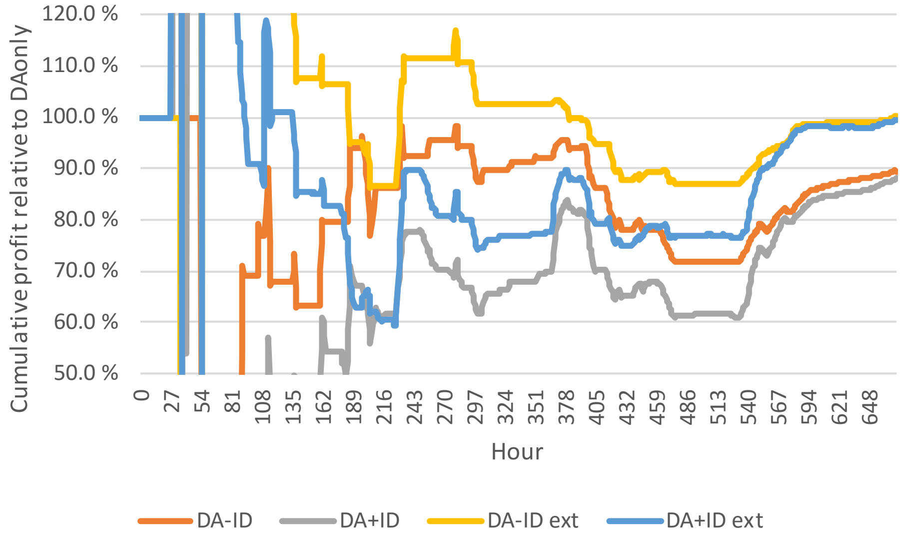

Figure 21 Cumulative profits with ID-trade relative to DAonly, first 4 weeks

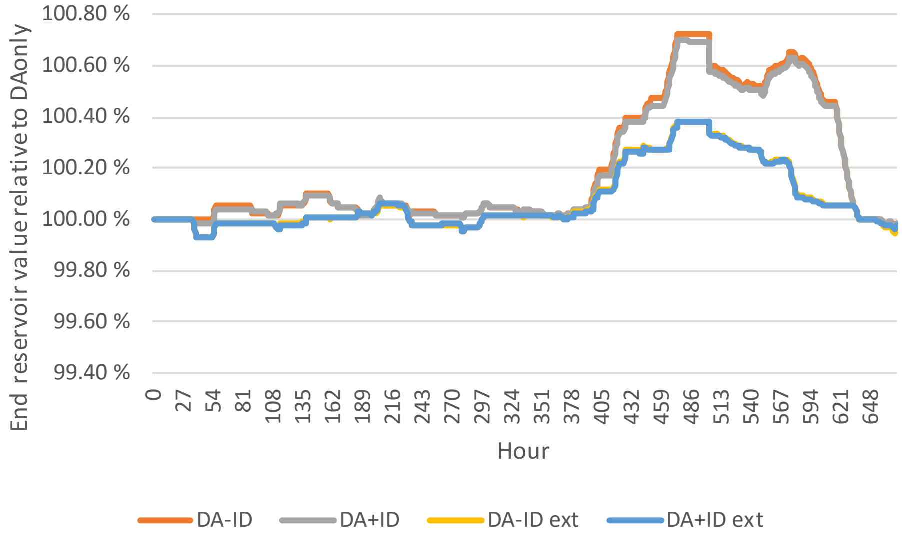

Figure 22 End reservoir values relative to DAonly, first 4 weeks

Though far less extreme, also the simulation period after week four, shows substantial variation in the ID trading performance. This is more easily seen in Figure 23 where the range of the vertical axis is narrower. An important consequence from this is that the time interval selected for assessment can strongly affect the conclusion on which trading strategy is the best. For instance, if the cumulative net profits are calculated and compared from the end of the inflow spike in week 4, where all simulations have the same state with full reservoirs and both generators spinning, the conclusion will be a small gain from ID trade relative to DAonly, as shown in

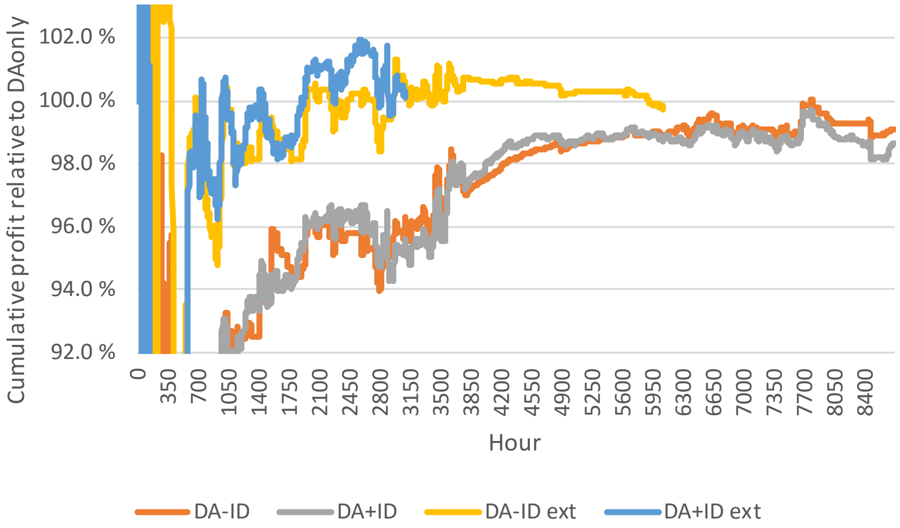

Figure 23 Cumulative profits relative to DAonly, zoomed vertical axis

Table 4 Cumulative net profits relative to DAonly for different simulation intervals using the original subproblem structure

	DA-ID	DA+ID
Whole year	99.2%	98.9%
Year less week 1-4	100.6%	100.3%

6.3 Variation in hourly income

This subsection looks into how the hourly income is distributed with the different market combinations. Since the plant has inflow giving only 42% full-load hours it is not surprising that all simulations show a substantial number of hours without production and income. In the following plots, showing the distribution of hourly income normalized[6], the vertical axis cuts the no-trade decision to improve readability for the rest of the plot. Comparing the distribution for DAonly (Figure 24), DA-ID (Figure 25) and DA+ID (Figure 26) we observe a clear difference in distributions, where DAonly has one peak, DA+ID has two and DA-ID is in-between (ignoring the no-trade instances). More importantly, market combinations with ID show a larger spread. DAonly loss is limited by the start/stop cost and the income is limited by the production capacity and DA price. DA-ID and DA+ID can on the other hand buy ID and thereby incur lager losses, but also achieve larger gains as the maximum of both DA and ID price limits the income. Some main characteristics of the three income distributions, with numbers given relative to the mean DAonly income, is given in Table 5.

Table 5 Characteristic of income distributions from simulation with original subproblems.

	DAonly	DA-ID	DA+ID
Mean	100%	99.4%	99.1%
Standard deviation	134%	134%	141%
Min	-17%	-67%	-500%
Max	998%	1028%	1272%

[6] "Normalized" implies that the sum of all bars in the plot is 1.

The simulations with extended model horizon show similar results for the 3086 simulated hours, but the difference in simulation period makes the results not directly comparable. A difference is that the rightmost peak is less distinct in the DA+ID distribution with extended model horizon. Plots and numbers for these simulations are given in the appendix.

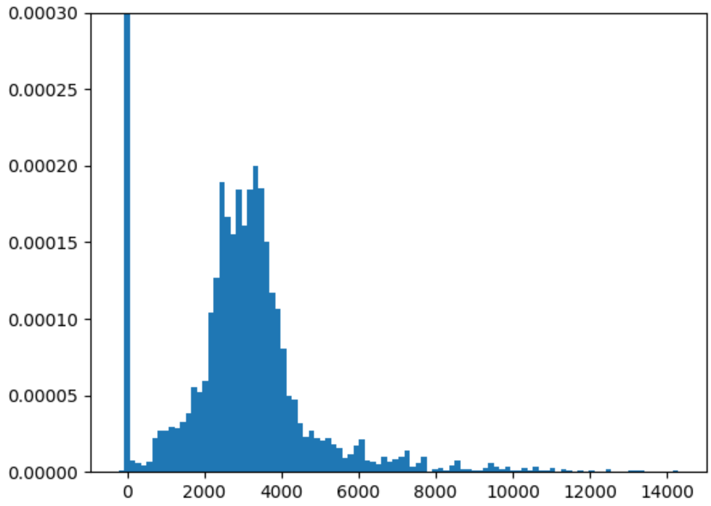

Figure 24 DAonly normalized hourly income distribution

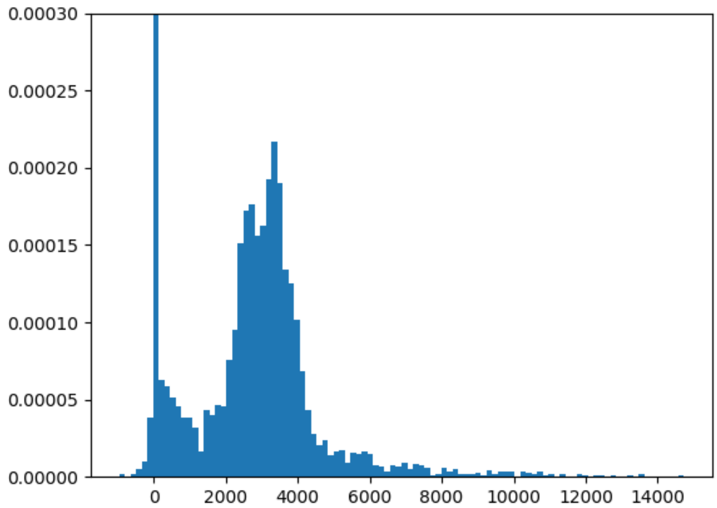

Figure 25 DA-ID normalized hourly income distribution from original simulation

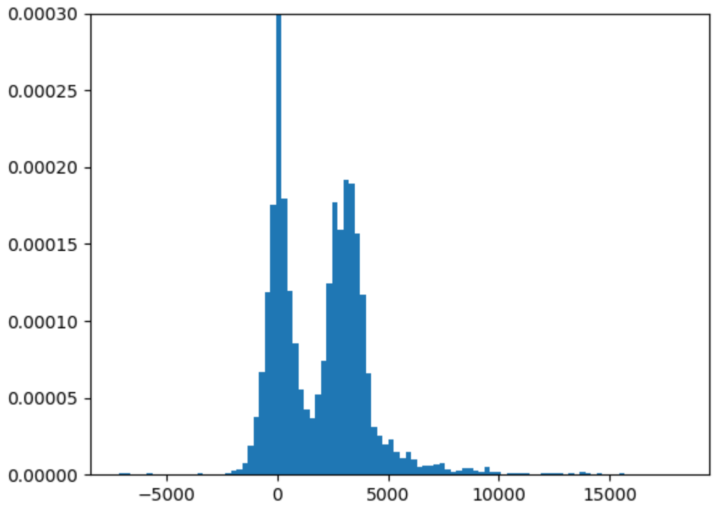

Figure 26 DA+ID normalized hourly income distribution from original simulation

6.3.1 Forecasted vs realized ID price

The inclusion of ID market trade both increases the action space, but also adds more short-term uncertainty to the simulations. When DA sales are decided at noon day-ahead, ID prices in the subproblems are uncertain. We calculate the expected ID price at the time of DA bidding, denoting this the ID forecast, and study two situations:

- "High": hours where the ID forecast exceeds the DA price while the ID realized price is below the DA price
- "Low": hour where the ID forecast is below the DA price while the ID realized price is above the DA price

In such situations the ID forecast gives misleading incentive to DA+ID when making DA decisions. In "High" the incentive is to withhold capacity from DA to be able to sell in ID (assuming both prices are above the water value). In "Low" the incentive is to sell max in DA to buy back in ID (assuming the ID price is below the water value) and thereby make a margin without consuming water.

Figure 27 and Figure 28 show the DA+ID normalized hourly income distribution over all hours ("all") and when hours with "High" and "Low" ID forecast is excluded ("excl High" and "excl Low"). Mean and minimum hourly income relative to mean hourly income with all simulated hours are reported in Table 6. The histogram clearly shows how a substantial share of the negative incomes are due to low forecasts causing unfavourable DA sale. Avoiding the high forecasts has a less visible effect on the income distribution. From the table, on the other hand, we observe that avoiding these high forecasts has a larger positive effect on the average income because a few large losses are avoided. The table also show that being able to predict if the ID premium[7] will be positive or negative would give zero or positive income in all hours.

[7] The ID premium is the price difference between the DA and the ID price

Table 6 DA+ID income characteristics when selected hours are removed

	all	excl High	excl Low	excl High and Low
Mean	100%	110%	105%	122%
Min	-505%	-505%	-27%	0%

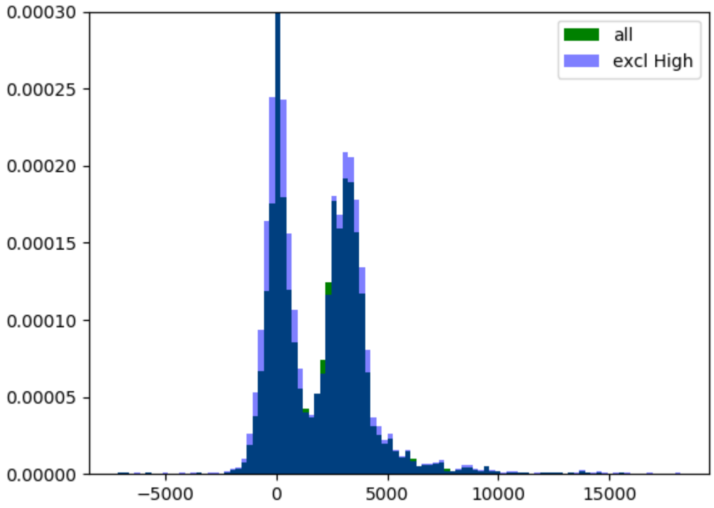

Figure 27 DA+ID normalized hourly income distribution from original simulation. Dark blue is the area where the two distributions are overlapping

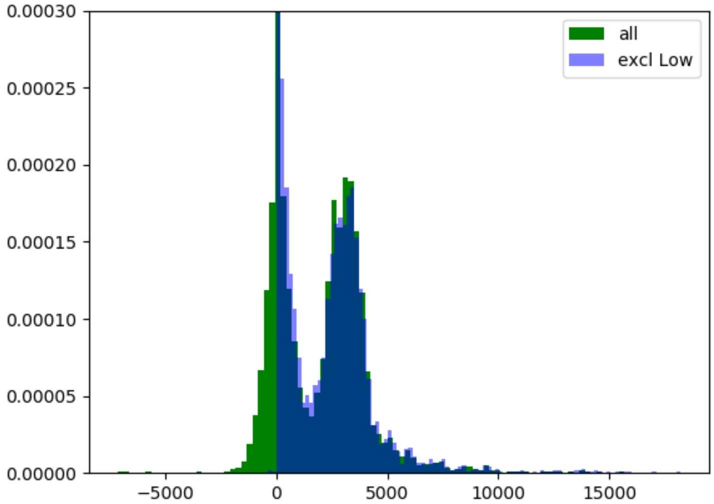

Figure 28 DA+ID normalized hourly income distribution from original simulation. Dark blue is the area where the two distributions are overlapping

6.4 Trade decisions

This section reports on market trading patterns for simulations with the original subproblem structure. The same patterns are observed also with extended subproblem horizon, but there are less hours with full production in this simulation. Plots corresponding to those given here, but based on results with extended subproblem horizon, are given in the appendix. With coordinated trading, in DA+ID, SHARM takes more extreme positions in the DA market relative to sequential trading. Actually, all positions are either at maximum production or zero, as can be seen in the left histogram in Figure 29. This indicates that despite the ID price uncertainty at the time of DA trading, the model does not find value in keeping a flexible position to allow deciding the direction of trade in ID when price uncertainty is reduced. It rather seems to take extreme positions based on the relative difference between DA and expected ID prices and use ID to adjust the final market position to a preferable production level. This agrees with the less extreme positions in the right histogram. The sequential model DA-ID use the ID market less actively, having far more zero-positions in ID, which is natural since the DA position in sequential trade is taken without a planned repositioning in ID.

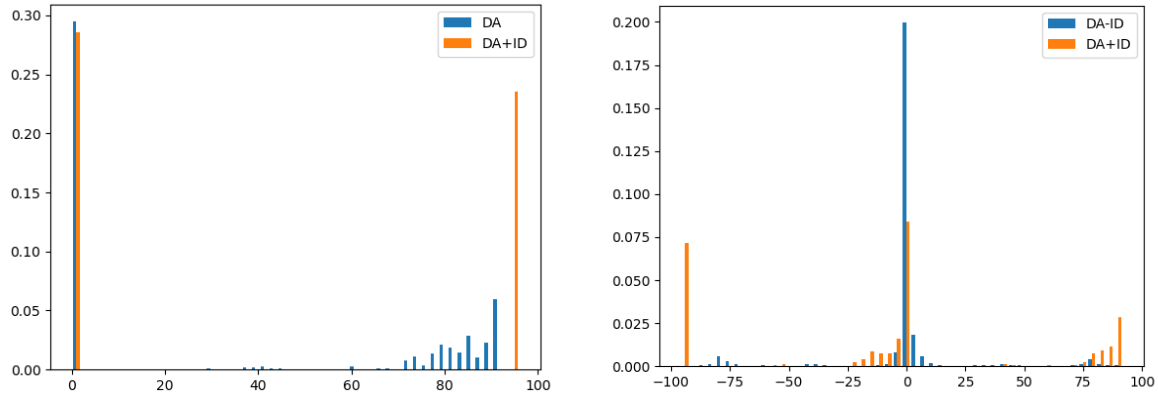

Figure 29 Normalized histograms over hourly net sales positions [MW] in DA (left) and ID (right)

The histogram over final trade positions in Figure 30 show that purely trading in DA implies less hours operating both at maximum and at zero output.

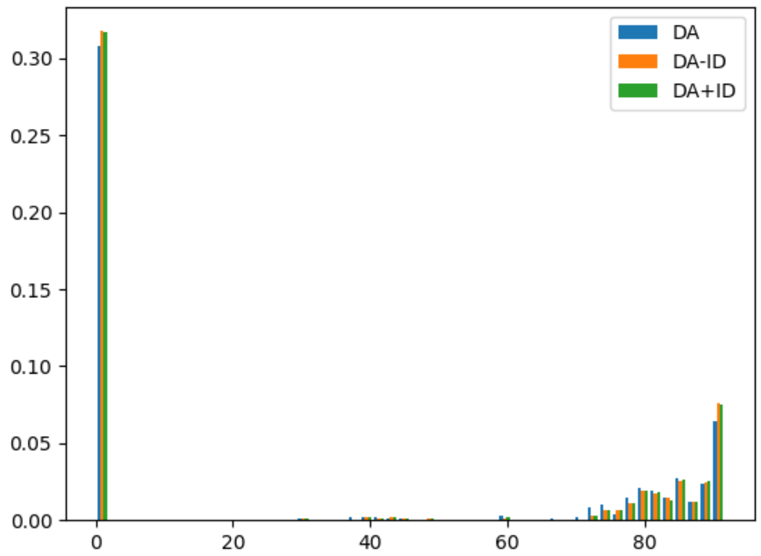

Figure 30 Normalized histogram for hourly total trade positions

6.5 Reservoir operation

The reservoir paths for simulations with original and extended subproblems are presented in Figure 31 and Figure 32, respectively. While the main patterns are equal across all simulations, DAonly has a tendency to keep a lower reservoir level, which is even clearer when looking at the duration curves in Figure 33 and Figure 34. The effect is present also when excluding the first four weeks covering the example from Subsection 6.1.1 and it is strengthened when extending the subproblem model horizon.

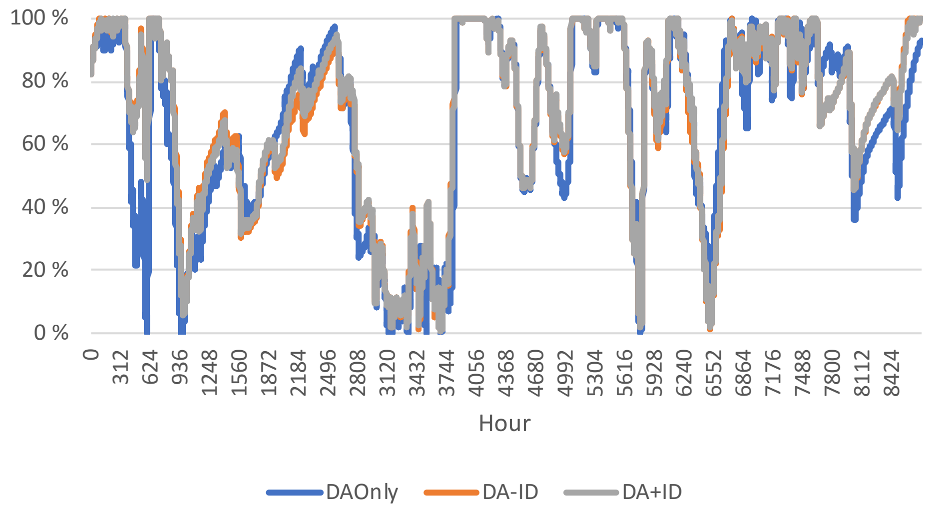

Figure 31 Reservoir paths from simulations with subproblems

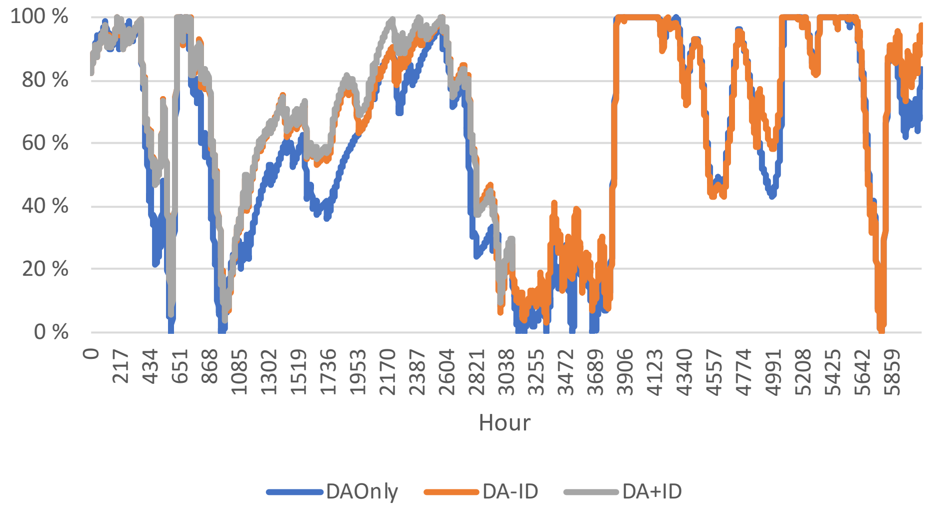

Figure 32 Reservoir path from simulations with subproblems with extended model horizon

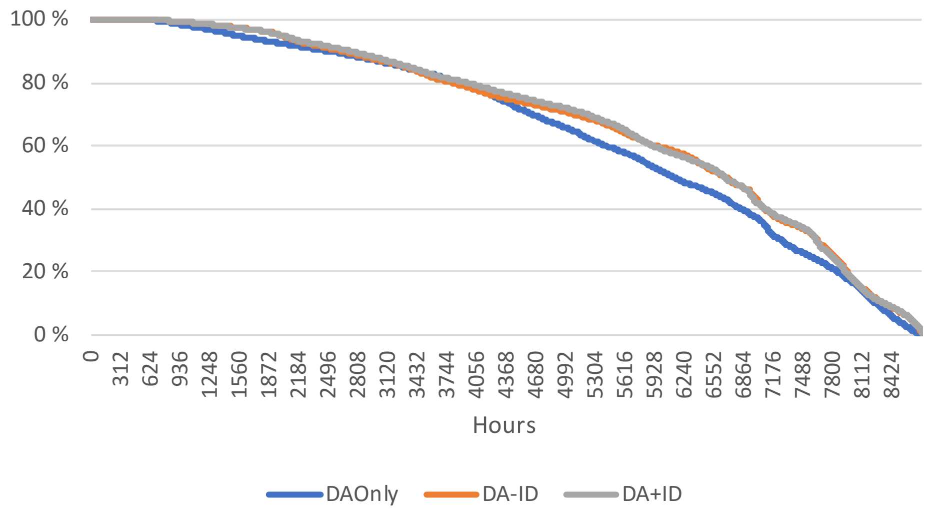

Figure 33 Duration curves for reservoir level from simulations with original subproblems

All simulations show some hours with very low reservoir levels which might seem unrealistic. To understand this, please remember that no inflow uncertainty is modelled in the short-term subproblems solved in SHARM. Still, the simulations including ID trade never go below 1% reservoir filling, as opposed to DAonly which goes below 1% 59 hours. A likely explanation for this difference in behaviour is that the subproblems containing ID trading only have one hour with deterministic prices as opposed to 24 deterministic hours in DAonly. Avoiding a completely empty reservoir leaves flexibility to utilize high ID prices if they show up while the 24 hour deterministic period makes it possible to run empty within this period and still rebuild some inventory before the next uncertain price.

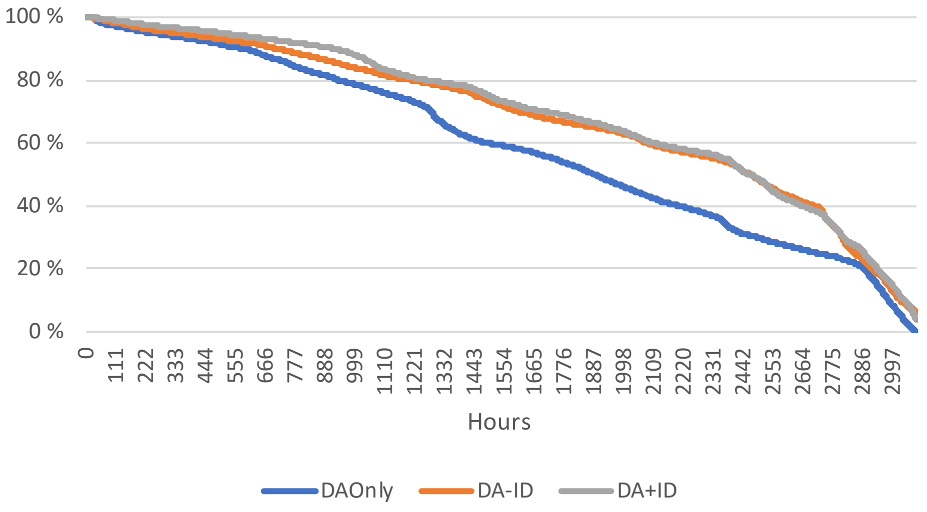

Figure 34 Duration curves for reservoir level from simulations with subproblems with extended model horizon

6.6 Start/stop

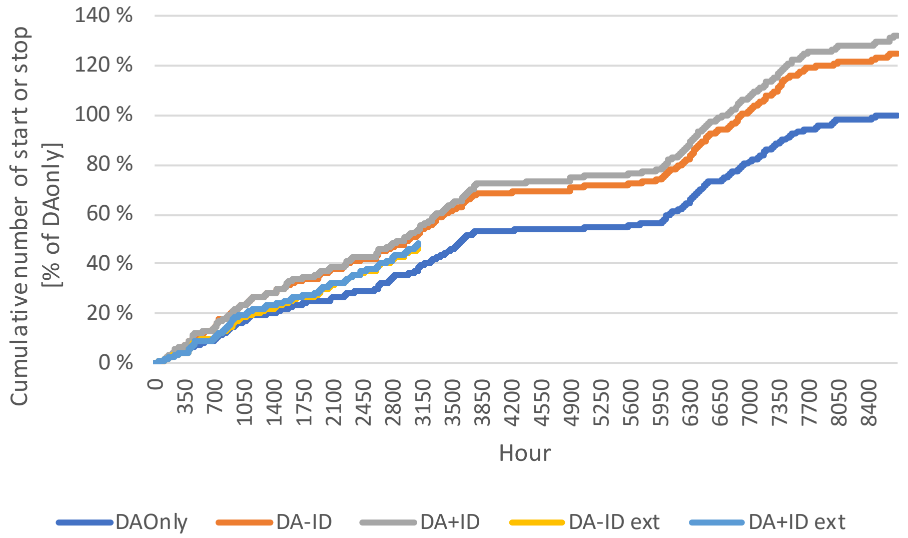

Figure 35 Cumulative start/stop costs

Both starting and stopping a generator trigger a start/stop cost. Figure 35 shows the cumulative number of start or stops relative to the total number for DAonly with the original subproblems. Not surprisingly due to the higher variability in ID prices relative to DA, allowing ID trade increases the number of starts and stops. This effect is still present but reduced when the subproblem model horizon is extended.

7 Final remarks

A main objective for this study was to build knowledge on the added value that can be achieved by trading in both day-ahead (DA) and intraday (ID) markets relative to DA sale only. Furthermore, we wanted to observe if coordinated trading (DA+ID) achieves a significantly larger value than sequential trading (DA-ID). The results indicated on average rather small, also negative changes in net profit with ID trading relative to DAonly, both for sequential and coordinated trading.

The comparison of cumulative net profits shows that especially DA+ID, but also DA-ID varies around the DAonly value. A consequence of this is that any attempt to measure the value of ID trade relative to DAonly will be sensitive to the time period over which the assessment is conducted. As earlier explained, in our study, a special situation within the first four-week period strongly affects the results. Looking at the original simulation over the whole year, we observe that neither DA-ID nor DA+ID reaches the cumulative profit from DAonly. When compensating for differences in end reservoir value the final results for DA-ID and DA+ID relative to DAonly are 99.2% and 98.8%, respectively. If we on the other hand measure the cumulative net profits from the end of the inflow spike in week four (when all simulations have equal and full reservoir), the respective values are 100.6% and 100.3%.

It should be noted that the results are measured towards an optimized DA-operation, DAonly, which naturally is a challenging benchmark. A pure ID trading strategy is not tested, so our results only assess the value of ID trade as a supplement to DA trade, not as an alternative. Furthermore, there can be multiple reasons for participating in the ID market. Unpreferable DA market clearing and operational uncertainties, i.e. inflow, can give a motivation for repositioning that is internally driven. On the contrary, the hydropower plant might have lower flexibility costs than other ID market participants, i.e. due to ramping and storage capabilities, making ID market participation profitable. As this study does not model neither operational uncertainty nor DA bidding, the market flexibility value is what is sought observed here.

This variability in cumulative net profits is to a large extend driven by the variability in income, as the start and stop cost are relatively small. The findings in Subsection 6.3.1 indicate that the match between the expected ID premium at the time of DA trading and the realized ID premium strongly affects the performance of DA+ID. This is decided by the price modelling, which in this study is purely driven by historical data and statistical methods. The individual power producers have own confidential procedures for price forecasting that most likely deviate from what is used here. Furthermore, from economic theory it can be questioned whether there should be an ID price premium at the time of DA trading or if ID price at this time stage on expectation should correspond to the DA price. Enforcing such a requirement on the ID price modelling is likely to make the model less eager to plan counter trading between DA and ID and thereby avoid losses due to forecast errors but also miss out on ex post profitable counter trading opportunities.

Moreover, the increased spread in hourly income when allowing ID trade relative to DAonly can be seen as an increased risk. Depending on company risk profile and risk management approach this might call for risk measures, for example conditional-value-at-risk (CVaR) or a limit on deviation from DAonly. Such measures have not been modelled in this study. In addition to reducing the hourly income spread, such measures could give the model an incentive to take more flexible DA positions in DA+ID rather than the extreme positions observed in Subsection 6.4.

The results show that how the end-of-horizon (EOH) effect is modelled in the simulation setup can strongly affect the simulation results. Our initial approach used rather short model horizons for subproblems with ID trading only, based on the idea that this should be hour-by-hour within-day repositioning, while the DA trade should capture the longer time perspectives. As explained in the subsection on model horizon, Subsection 6.1, this approach did give unpreferable behaviour in the ID market. The following simulation results with and without the extended subproblem horizon confirm that the EOH modelling is important not only for the daily to weekly allocation of water in DA, but also for the hourly ID decisions. In our study, we improved the EOH modelling by increasing the model horizon of the subproblems. This is one way of doing it that improved comparability between the conditions for DA and ID decisions in our simulations. Other options could be to limit the decision space for ID trading by e.g. using a reservoir schedule limiting how much ID trade is allowed to change the reservoir path relative to the DAonly reservoir path in the end of each subproblem. This limits unpreferable, but also possibly preferable deviations from the DAonly operation. Another approach would be to improve the water value description. In our setup we used water value functions (cuts) with a weekly time granularity, which is rather coarse. The improvements observed through increasing the subproblem model horizon indicate a significant difference between the value of water implicitly described in the extended model horizon compared to the value given by these weekly water functions. Some possible changes that could improve the water value functions could be to increase the time granularity to e.g. daily water value functions to give a smoother description rather than clear shifts each time the simulation horizon reaches a new water value function. This would probably make comparison of different model runs (like DAonly and DA-ID/DA+ID) less sensitive to differences in when each model ends. Furthermore, using water value functions from SHARM rather than ProdRisk could give a closer match between the model's implicit representation of water value and the one given by the water value function. In our simulation setup, reservoir paths from DAonly could be an indicator of relevant reservoir levels which could limit the number of model runs necessary to calculate water values spanning a sufficient reservoir range. Further studies are necessary to assess the performance, both in result quality and computational burden, for these options.

A previous study (Fodstad, Aarlott, & Midthun, 2018) evaluated the value of multimarket hydropower trade with perfect foresight. Data from the German market from 2015 was used, and the added value compared with DA trade only was calculated for different plant characteristics. As this study used data from 2017 the results are not directly comparable, but still represent a relevant reference level. As the modelled plant in this study is relatively inflexible, the Run-of-river plant from (Fodstad, Aarlott, & Midthun, 2018) is the most closely comparable. Full load hours are 48% and 42% for Run-of-river and the current respectively, while degree of regulation is 0 and 0.017 respectively. (Fodstad, Aarlott, & Midthun, 2018) observes 1.1%[8] added value when allowing perfectly coordinated ID trade in both directions compared to DA sale only. Similar numbers for a more flexible plant "Small", with 48% full load hours and degree of regulation at 0.75, is reported at 1.7%. Our simulations do not confirm similar gains from ID trading on top of DA trading. As (Fodstad, Aarlott, & Midthun, 2018) assumed perfect foresight in all prices, the difference in observed gains might indicate that the modelling of end-of-horizon and prices (forecast and scenario tree generation) are decisive for the ability to realize added value in ID.

A Supplementing results

A.1 Variation in hourly income

Table 7 Characteristic of income distributions from simulation with extended subproblem horizon

	DAonly	DA-ID	DA+ID
Mean	100%	101.1%	100.9%
Standard deviation	228%	231%	251%
Min	0%	-64%	-761%
Max	1517%	1533%	1934%

[8] This and similar numbers are calculated based on result raw data presented in the publication.

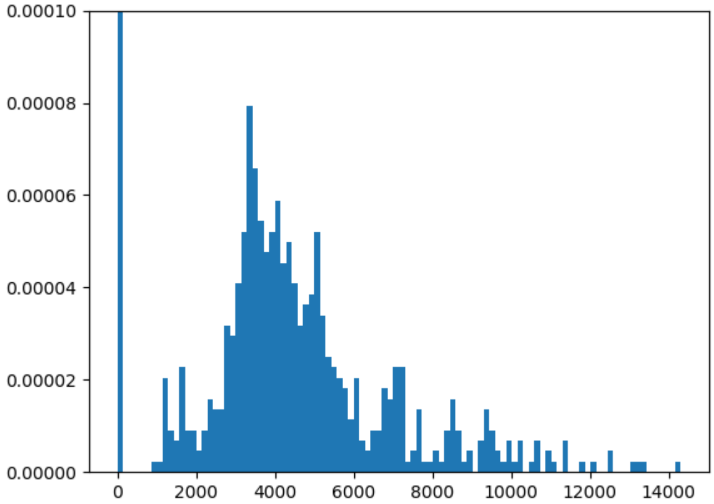

Figure 36 DAonly normalized hourly income distribution for the first 3086 hours

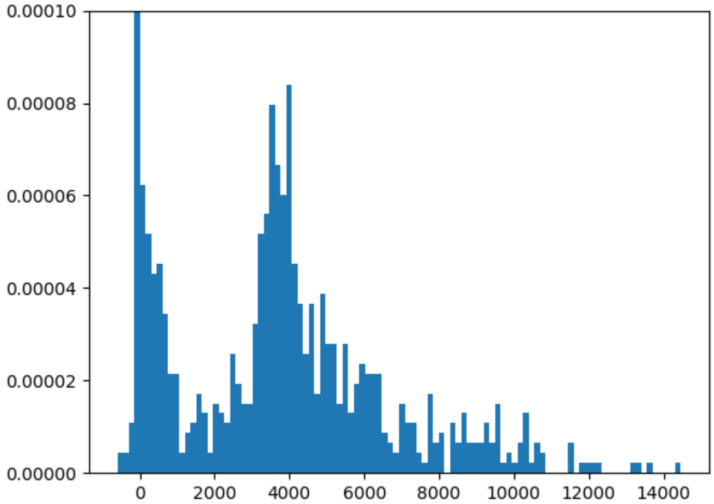

Figure 37 DA-ID normalized hourly income distribution from simulation with extended subproblem horizon

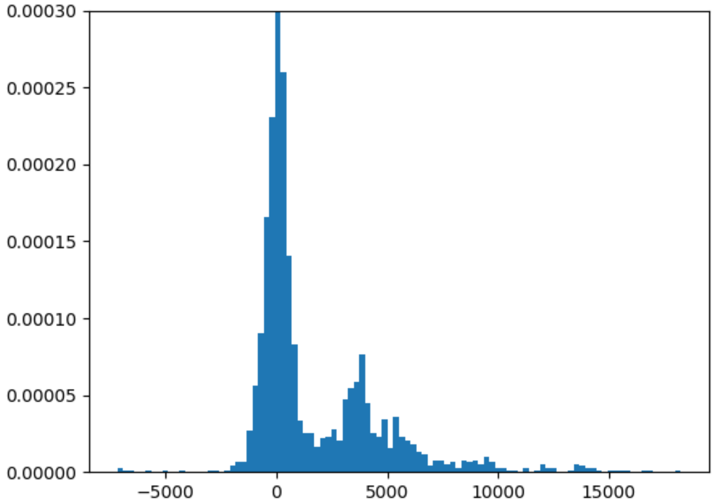

Figure 38 DA+ID normalized hourly income distribution from simulation with extended subproblem horizon

A.2 Trade decisions

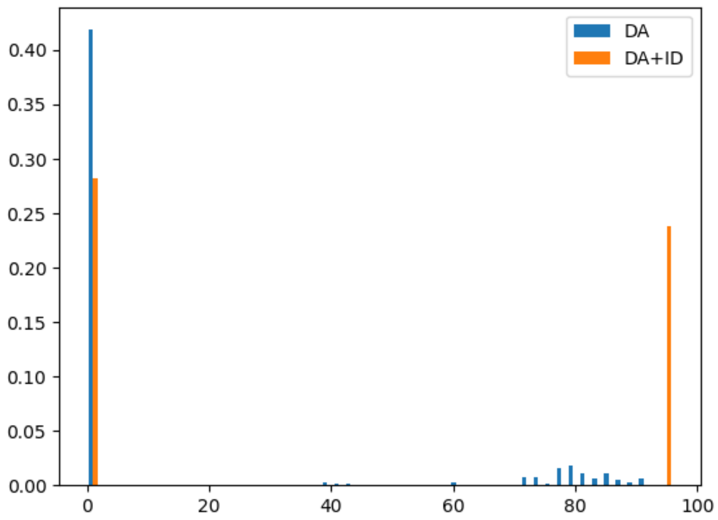

Figure 39 Normalized histograms over hourly net DA sales positions [MW] from simulations over 3086 hours with extended subproblem horizon

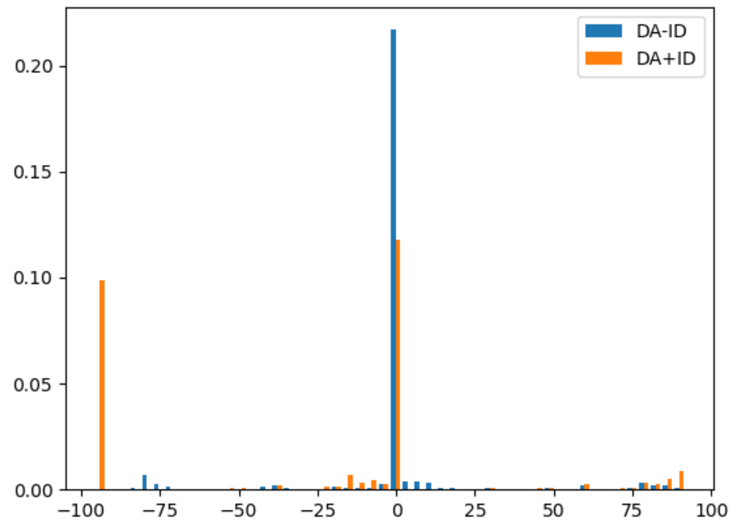

Figure 40 Normalized histograms over hourly net ID sales positions [MW] from simulations over 3086 hours with extended subproblem horizon

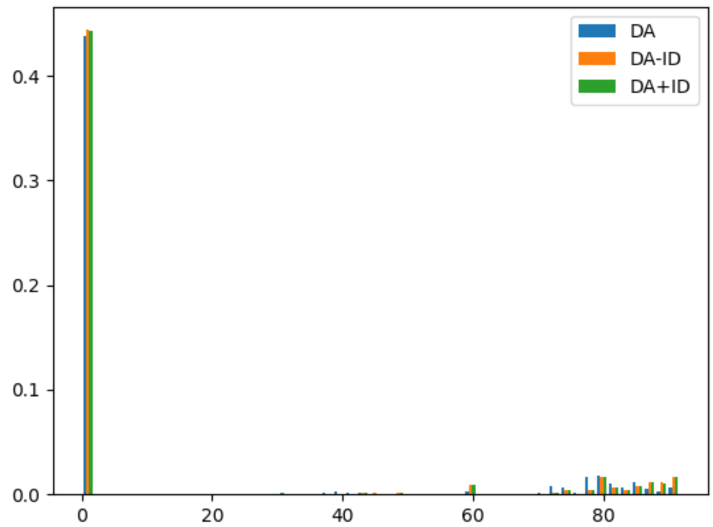

Figure 41 Normalized histogram for hourly total trade positions from simulations over 3086 hours with extended subproblem horizon

SINTEF

Technology for a better society

www.sintef.no