The value of coordinated hydropower bidding in the Nordic day-ahead and balancing market

Ellen Krohn Aasgård[1]

Abstract

Modern power markets often consist of a series of sequential markets where power is traded. Market agents must coordinate their trading strategy across all markets. In this paper, we formulate the mult-market bidding problem for power producers as a stochastic program. This formulation is implemented within the framework of optimization models used by the Nordic hydropower industry. We also present how input to the stochastic program may be generated by using a forecast-based scenario generation method combined with time-series models that predicts future prices and turnovers in the markets. The model is applied in a case study to investigate the value of participating in the Nordic day-ahead and balancing market. A producer may participate in the balancing market either by considering the markets sequentially or coordinated. Using cases with limited and perfect information about the balancing market to calculate lower and upper bounds, we find that the value of participating in the balancing market using the sequential approach is between 0.8 and 2.6%. Using the coordinated approach, the producer may gain between 1.4 and 2.9%. The value of coordination, i.e. the value of using the coordinated over the sequential approach, is found to be higher in the limited information case (1.7%) than in the perfect information case (1.1%). This indicates that, the more uncertainty, the higher is the value of coordination.

1 Introduction

Modern, deregulated power markets often consist of a series of markets where power is traded. Electricity may thus be traded several times before actual consumption or production takes place. The fact that there are several markets for the same product gives rise to a complex decision situation for market agents. In addition, the

---

✉ Ellen Krohn Aasgård
Ellen.Aasgard@ntnu.no

1 Norwegian University of Science and Technology, Alfred Getz vei 3, 7491 Trondheim, Norway

Springer

various markets represent different degrees of flexibility. In the day-ahead market, power may be traded well ahead of the operating hour, which may be necessary for inflexible technologies with long warm-up times or start-up costs. In the intraday and balancing market, power may be traded closer to real-time, which may be necessary for technologies that are more difficult to schedule in advance, such as wind power. For flexible technologies such as hydropower, participating in the intraday or balancing market may bring increased profits.

Coordinating trades in sequential power markets is a complex challenge that combines market aspects with details of the physical production system. This paper presents how coordination of trades across multiple markets may be formulated within the framework of the optimization models used by the Nordic hydropower industry [1]. We also show how a set of time series models that describes the nature of prices and turnovers in the Nordic power markets are used for scenario generation to create input to the optimization model. Using the optimization model and the scenarios, we analyze the value of participating in and coordinating trades across multiple markets. By letting the market agent have limited or perfect information about the balancing market, we develop lower and upper bounds on the value of coordination. The main contribution of this paper is to show how these bounds are calculated for a case study representing a realistic hydropower system.

In [2] the multi-market bidding problem is defined as the problem faced by a producer that participates in several markets. A producer may take two approaches to the multi-market bidding problem: sequential or coordinated bidding. Sequential bidding is when the producer participates in both the day-ahead and an additional market, but consider the two markets sequentially and determines the volumes in the day-ahead market without considering the other market(s). In the coordinated case, the producer decides its trades in the day-ahead market while accounting for potential trades in the other market(s). In this way, the trading strategy is optimized across markets as a joint problem.

A joint optimization problem requires more computational resources than sequential optimization. It is therefore of interest to determine the value of coordination. If the value of coordination is large, it is well worth the effort to model this more complex problem. However, if the value is small, producers might be better off using the sequential approach and rather improve other parts of their operations or modelling. In fact, very few Nordic hydropower companies use formal optimization methods to determine the trades in markets beyond the day-ahead market. The day-ahead market is the main market place for power and therefore receives the most attention in the planning process. Trades in the other markets may be done more on an ad-hoc basis. The value of both the sequential and coordinated approach may therefore be compared to participating in the day-ahead market only, in order to find the value of participating in the other markets. Bounds on the potential added value for market agents that participate in the Nordic day-ahead and balancing market are calculated in [2]. They find a gain of participating in the balancing market that varies from 9.60 to 30.79% over the months of the year. The gain of coordinated versus sequential bidding was found to vary between 8.03 and 24.79% over months of the year. Additional tests show that the gain of coordinated bidding do not necessarily increase with an increase

in balancing prices (as this increases the value of up regulation but decreases the value of down regulation) but rather with an increase in the difference between balancing and day-ahead prices.

A complicating factor when comparing literature on multi-market bidding, is that market arrangements and terminology vary around the world. Most power markets consist of a series of sequential markets where power and related products are traded. The markets differ in respect to the product traded (energy, capacity or ancillary services), the time to delivery, the auction or market clearing format, and demand side competition. In Europe and most North American regions, the wholesale power market consists of a day-ahead market, where expected production and consumption for the next day is traded, followed by an intraday market that may have continuous trade and a real-time balancing market for handling instantaneous imbalances in the power system.

The day-ahead market works pretty much in the same manner around the world. However, possibilities for adjustments in intraday markets and procurement of reserves for the transmission system operator (TSO) may take many forms depending on institutions, energy mix and traditions in each area. In this paper we use the rules and framework of the current Nordic market operated by NordPool, but the methods presented may be applicable to other market arrangements with some modifications. Specifically, the Nordic market consists of a day-ahead market, an intraday market with continuous trade and a real-time balancing market where the TSO procures reserves. We consider joint optimization of trades in the Nordic day-ahead and balancing market in this work, as this was previously analyzed by [2].

In existing literature, a common approach for the multi-market bidding problem is stochastic programming [3], which has been applied by [2, 4–8]. Stochastic programming is also used in this work. Formulations based on stochastic programming require that the stochastic parameters. i.e., the prices and turnovers in the markets, are described by scenario trees. Several methods have been suggested for generating scenario trees [9, 10], each using different methods and requiring different types of information about the stochastic parameters. Sometimes the best available information about an uncertain future is a single forecast. A single forecast does not provide enough information to construct a scenario tree, but may be combined with historical data on forecast errors to produce a scenario tree using the method in [9]. In the current work, time-series models for the prices in the Nordic day-ahead and balancing market are applied to generate forecasts that are used together with the forecast-based scenario generation method to generate scenarios trees.

The rest of this paper is organized as follows. In Sect. 2, we formulate the multi-market bidding problem for hydropower as a stochastic problem. In Sect. 3, we describe how two existing time-series models, namely the models developed by [13, 14], have been fitted to data for prices and turnovers in the day-ahead and balancing market. The resulting forecasts are used together with the forecast-based scenario-generation method [9] to generate scenarios trees as described in Sect. 4. Sect. 5 contains results from a case study where the optimization model is applied to an example hydropower system. Section 6 gives some conclusions.

2 The multi-market bidding problem

This section presents the basic mathematical modeling of a stochastic optimization model that coordinates multi-market trades for a power producer. Each market is described, in every time-step t and scenario s, by a price (or a premium) and a limit on the maximum volume that may be traded at that price. The volume limit is used to describe markets where there is a limited turnover, such as the Nordic intraday or balancing market. The volume limit, V_{ts}^m is a time series and may be a stochastic input to the model. In every time step and scenario, the volume y_{ts}^m sold in any market m must be less than the demand in the market, i.e.

$$y_{ts}^m \leq V_{ts}^m. \tag{1}$$

Prices, P_{ts}^m, may naturally also be stochastic. The objective function is to maximize profits from sales in all markets, as in

$$max \sum_t \sum_s \sum_m \pi_s \left(P_{ts}^m y_{ts}^m \right), \tag{2}$$

where the summation is over all time-steps, all scenarios and all markets. π_s is the probability of each scenario and y_{ts}^m is the volume offered in each market. y_{ts}^m must be subtracted if the producer buys energy from the market (e.g. for pumping or for substituting own production if cheaper energy is available) and when the producer offers downward balancing, otherwise the sales in each market are added. The sum of the position made in each market must be equal to the produced volume (perhaps with some small deviations $z_{ts}^{+/-}$), according to

$$\sum_m y_{ts}^m = \sum_g x_{tsg} + z_{ts}^+ - z_{ts}^- \tag{3}$$

where x_{tsg} is the produced volume from each generator g in time step t and scenario s. The deviations $z_{ts}^{+/-}$ are penalized in the objective function in order to reflect the cost of trading these unbalances in yet another (imaginary or real) adjustment market. The above model assumes that any volume x_{tsg} may be produced. Actual production systems are much more complex. Details of hydropower production are however omitted from the presentation here. In fact, the above model may be used by any producer that participates in multiple markets as long as it is combined with a representation of the specific production system. In our case, the multi-market model is implemented within the framework of optimization models that is used for short-term production scheduling by most large hydropower producers in the Nordic region [1]. The volume to be produced from the generating units, x_{tsg}, is thus determined by this more complex model that includes all technical, hydrological and environmental constraints relevant for hydropower production, e.g., minimum production levels, forbidden operating zones, start/stop, discharge dependent losses in tunnels and penstocks, minimum and maximum reservoir levels, minimum and maximum river or tunnel flows and more. This also includes more terms in the objective function, for instance start or stop costs that are modelled by integer

variables. A value of water left in the reservoir at the end of the optimization period is also included. This may be interpreted as the marginal cost of production and is referred to as the water value in hydropower scheduling. A description of the production scheduling model can be found in [11].

The simple model formulation above is however not complete without modelling the information structure in the multi-market setting. Scenario trees should represent how information is revealed to the market agent over time. In the case of a market agent that coordinates its trades across sequential markets, the prices in the different markets will be revealed in sequence over time. The day-ahead prices are revealed once every day when the market clears, while the balancing market prices are revealed in real-time. The decisions for each market must also be taken in sequence. In the current set-up of the Nordic power market, the trades (or bids) for the day-ahead market must be submitted at noon the day before operations. Trades in the balancing market may be submitted up to 45 min before the delivery hour. This market structure must be reflected decision problem. Particularly, decisions in all markets must be made prior to the prices for that market are revealed. In terms of a stochastic programming problem, the trade decisions cannot depend on any particular scenario being realized. The question is then to maximize profits across all markets given the information that is available to market agents at any stage in the decisions process. The need for coordination across markets arises because the final commitment, i.e. the actual volume to be produced in a specific hour, is the summation over the position made in each market. For every hour, agents must make sure that their trades summarize to a feasible output level.

Consider the left part of Fig. 1 for a closer look at the scenario tree structure for the multi-market problem. The day-ahead prices are revealed once every day when the market clears. This means that the scenario tree for day-ahead prices must have a new stage every 24 h. For a 72-h horizon where the scenario tree branches into two
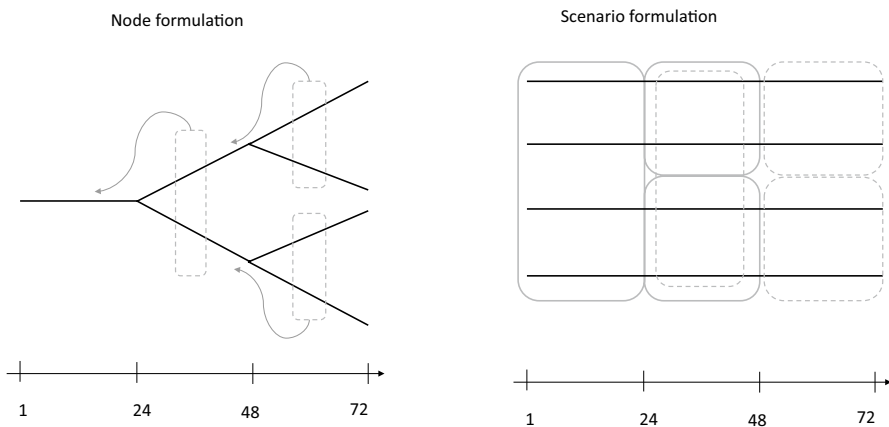

Fig. 1 (Left) Node representation of a scenario tree with daily branching. The dotted grey boxes illustrate that the traded volumes must be determined before prices are revealed. (Right) Scenario representation of a scenario tree with daily branching. The fully drawn grey boxes illustrate the normal non-anticipativity constraints, while the dotted grey boxes represent the market non-anticipativity constraints

new scenarios at each branching step, daily branching would yield $2^2 = 4$ scenarios. In the balancing market, however, prices and volumes are revealed in real-time, which would lead to a scenario tree with hourly or even finer branching. This would quickly lead to a very large problem, especially considering that more than two new scenarios at each branching point is necessary to represent the full uncertainty of prices. To avoid this curse of dimensionality, we choose to have daily branching also for the balancing market, i.e., that both the day-ahead prices and the balancing market prices and volumes are revealed together when the day-ahead market clears. This assumption means that the models sees no uncertainty in the balancing market during each day.

In the stochastic program, we use a scenario representation rather than a node formulation, see the right part of Fig. 1. This means that we must explicitly include non-anticipativity constraints stating that if two scenarios s and s' are indistinguishable at time t on the basis of information available at time t, then the decisions made in scenario s must be equal to the decisions made in scenario s'. In our case, this means that the traded volume must be equal between all scenarios belonging to the same node,

$$y_{ts}^m = y_{ts'}^m. \tag{4}$$

The non-anticipativity constraints are illustrated by the fully drawn grey boxes in Fig. 1. The optimization model also needs to know that day-ahead trades must be done prior to market clearing and that day-ahead trades must be made prior to balancing market trades. This means that the day-ahead trades cannot depend on any particular realization of prices for the next day. We call this the market non-anticipativity constraints and formulate them as

$$y_{ts}^{DAM} = y_{ts'}^{DAM} \tag{5}$$

where the superscript *DAM* denotes the day-ahead market. The market non-anticipativity constraints are illustrated by the dotted grey boxes in Fig. 2. Similar constraints may be applied to any market trade volume variable, y_{ts}^m, depending on whether the trades in the particular market are to be decided in real-time or not. If the constraint is imposed on the balancing trades and the current scenario tree structure is kept, it means that the balancing market trades must be determined prior to clearing of the day-ahead market, i.e. at the same time at determining the day-ahead market trades. If the constraints are not imposed, the balancing trades may be determined after clearing of the day-ahead market. Since there is no uncertainty within each day regarding the balancing market, this means that the producer has knowledge of the realized balancing market prices and volumes. This is of course not possible in reality, but we include it in our case study to measure the value of having perfect information in the balancing market. The results in the case study in Sect. 4 may thus be interpreted as bounds on the value of participating in the balancing market. If the producer has to decide the balancing market trades before clearing of the day-ahead market, this will underestimate the true value. If the producer can decide the balancing trades in real-time, i.e. with perfect information, this overestimates the true value. As stated earlier, daily branching of the scenario tree is used in

this work for computational reasons. If we had used hourly branching or blocks of hours being revealed during each day for the balancing market, we would have gotten better estimates of the value of participating in the balancing market.

3 Modelling the markets

In this work, we consider the Nordic day-ahead and balancing market. The day-ahead market is operated by NordPool and cover the Nordic and Baltic countries. The day-ahead market is the main market place for power in the Nordics with 394 TWh traded in 2017 [12]. The balancing market is the Norwegian market for tertiary reserve, also called the regulating power market or manual frequency restoration reserves (FRR-M) market. For both the day-ahead market and the balancing market, we use data for the Norwegian price zone NO2.

For each market, we fit time-series models that describe the main properties of the market, i.e. the price or premium and the turnover. The time-series models are used to forecast or simulate future values to be used for scenario generation. Most effort is devoted to developing a good model for the balancing market. To the author's knowledge, hydropower producers often have their own forecasting models for the day-ahead market, either in-house or supplied by external forecasting companies. The aim of this work is not to compete with such models. For the balancing market, however, good forecasts are harder to come by. Reasons for this is discussed in [13], which also formulate the balancing market model used in this work. The model for the day-ahead market is based on [14].

3.1 The day-ahead market

Due to the daily clearing of the day-ahead market, the hourly day-ahead market prices cannot really be represented as a pure time-series process. In normal time series, the information set is assumed to be updated when moving from one time step to the next. This is not the case for day-ahead prices, because the information set is updated on a daily rather than an hourly basis. The prices for all 24 h of the next day are set simultaneously in the market clearing. Thus, it is more correct to model the day-ahead prices as a time series of 24-h panel data rather than a single time series. Our method is based on [14], but we also account for seasonal variations. The model in [15] is a more advanced model in that it includes exogenous variables such as reservoir filling, temperature and wind power production forecasts, but we aim to keep the modelling more simplistic. Following [14], the day-ahead price P_{td}^{DAM} in hour t on day d is the sum of two independent components: a deterministic component f_{td} and a stochastic component x_{td},

$$P_{td}^{DAM} = f_{td} + x_{td}. \tag{6}$$

The deterministic component accounts for predictable regularities, such as the mean price level and daily and seasonal variations,

$$f_{td} = \mu_t + \sum_d \beta_t^d I_t^d + \sum_s \beta_t^s I_t^s. \tag{7}$$

The mean price level is denoted μ_t and may vary between hours. Daily variations are modelled by dummy variables I_t^d that equals 1 if the delivery day d is a specific weekday, $d = 1$ is Saturday, $d=2$ is Sunday ...and $d =7$ is Friday. The coefficient β_t^d is the daily and hourly difference from the mean level, and β_t^{Fri} is taken to be 0 to avoid collinearity. We also include dummy variables I_t^s for the season of the year. We choose to work with seasons instead of months or weeks to keep the model small. We define Winter to be weeks 48–8, Spring to be weeks 9–21, Summer to be weeks 22–34 and Fall to be weeks 35–47. Again, β_t^{Fall} is taken to be 0 to avoid collinearity.

The stochastic component accounts for the variation in price around the deterministic component. In this work, the stochastic component is an AR(1)-model, i.e. it accounts for the dependency between the prices in 1 h of today, p_{td}^{DAM}, and the same hour the day before, p_{td-1}^{DAM}. Notice that the AR(1)-process is not between consecutive hours on the same day (as in a normal time-series), but between similar hours on consecutive days. This is due to the daily clearing of the day-ahead market. The stochastic component thus becomes

$$x_{td} = \alpha_h x_{td-1} + \epsilon_{td} \tag{8}$$

where the errors ϵ_{td} will be *IID* $\sim (0, \Sigma)$.

To estimate the model, i.e. to determine the parameters μ_t, α_t, β_t^d, β_t^s and Σ, we use the method of seemingly unrelated regressors (SUR). SUR is appropriate for linear regression models that consists of several regression equations, each having its own dependent variable and potentially different sets of explanatory variables. In the case of day-ahead prices, we have one regression equation for each hour in the panel, that is, the system of equations is 24 copies of Eq. (6). The equation for each hour is a valid linear regression on its own and can be estimated separately. However, the error term may be correlated across equations, and it is likely that it is, since the bids for each hour underlying the market clearing prices are all quoted based on the same information set the day before [14, 15]. The correlation between error terms will be described by the covariance matrix Σ. The panel data model is estimated using SUR-techniques in the R package systemfit [16]. The fitted values can be found in Tables 1 and 2. Another important output from the SUR-estimation is the covariance matrix that describes the covariance between the error-terms in the stochastic component, Σ. The covariance matris for our fitted model is given in Table 3. Almost all hours have significant correlations with the other hours. This is different from [14] where mainly the peak-hours had significant correlations.

In the formulation of the optimization model in Sect. 2, each market must be modelled by a volume limit as well as a price. However, we do not model a time-series for the day-ahead market volume. The volume in the day-ahead market is never binding because the market turnover is much larger than the production capacity of an individual market agent. The volume limit in the day-ahead market is therefore taken to be a large, non-stochastic parameter.

Table 1 Estimated coefficients for the model for day-ahead prices. Hours 1–12

Effect	Variable	Hour											
		1	2	3	4	5	6	7	8	9	10	11	12
Average	$\widehat{\mu}_h$	64.37	79.46	90.63	94.61	100.80	100.24	101.73	151.43	174.01	167.62	154.51	150.11
Saturday	$\widehat{\beta}_h^{Sat}$	0.94	− 0.23	− 1.70	− 2.84	− 6.39	− 10.74	− 15.36	− 25.29	− 31.57	− 25.13	− 18.87	− 15.72
Sunday	$\widehat{\beta}_h^{Sun}$	− 6.96	− 9.17	− 13.96	− 12.08	− 14.70	− 19.26	− 21.44	− 28.59	− 32.80	− 26.09	− 19.35	− 15.40
Monday	$\widehat{\beta}_h^{Mon}$	− 1.18	− 2.34	− 0.54	− 1.05	3.25	13.28	22.04	21.63	22.51	19.56	16.29	14.31
Tuesday	$\widehat{\beta}_h^{Tues}$	1.21	1.33	1.11	1.71	2.79	2.42	4.19	5.03	2.72	1.76	2.20	2.96
Wednesday	$\widehat{\beta}_h^{Wed}$	− 4.01	− 5.07	− 6.25	− 6.54	− 5.36	− 2.28	1.23	3.77	3.57	− 0.24	1.17	1.45
Thursday	$\widehat{\beta}_h^{Thurs}$	− 2.56	− 3.81	− 3.79	− 2.52	− 0.99	0.25	2.85	3.33	3.94	1.48	1.00	0.79
Winter	$\widehat{\beta}_h^{Winter}$	− 0.47	− 0.13	0.45	1.05	1.01	0.32	0.53	9.01	19.44	16.43	12.55	10.58
Spring	$\widehat{\beta}_h^{Spring}$	− 3.43	− 4.45	− 5.60	− 5.02	− 5.98	− 6.29	− 5.88	− 4.19	− 1.03	− 3.66	− 5.75	− 7.17
Summer	$\widehat{\beta}_h^{Summer}$	− 12.71	− 16.73	− 20.49	− 22.79	− 26.27	− 25.25	− 24.05	− 34.53	− 38.25	− 35.63	− 32.15	− 31.33
Regression	$\widehat{\alpha}_h$	0.73	0.66	0.60	0.58	0.56	0.57	0.58	0.40	0.33	0.35	0.39	0.40

Table 2 Estimated coefficients for the model for day-ahead prices. Hours 12–24

Effect	Variable	Hour											
		13	14	15	16	17	18	19	20	21	22	23	24
Average	$\widehat{\mu}_h$	136.81	130.85	125.13	124.53	146.51	147.89	129.96	92.67	66.73	55.24	52.81	51.08
Saturday	$\widehat{\beta_h^{Sat}}$	−14.83	−14.63	−14.98	−11.35	−11.17	−9.57	−4.92	−4.49	−3.02	−2.71	−3.72	−4.94
Sunday	$\widehat{\beta_h^{Sun}}$	−13.31	−11.96	−11.45	−8.99	−9.38	−8.18	−2.09	2.68	5.97	6.80	5.02	1.64
Monday	$\widehat{\beta_h^{Mon}}$	14.66	16.59	17.88	22.15	23.76	23.46	18.32	12.08	9.62	7.96	4.22	1.54
Tuesday	$\widehat{\beta_h^{Tues}}$	3.89	5.44	5.09	5.88	3.97	3.30	5.61	3.37	4.47	4.39	2.66	−0.58
Wednesday	$\widehat{\beta_h^{Wed}}$	2.32	4.26	5.45	7.29	10.61	11.86	8.62	6.05	4.76	4.14	1.05	−1.76
Thursday	$\widehat{\beta_h^{Thurs}}$	0.71	2.17	2.11	4.34	2.43	1.16	3.78	2.25	3.03	2.95	0.68	0.03
Winter	$\widehat{\beta_h^{Winter}}$	8.24	7.98	7.62	9.56	15.95	20.32	12.72	3.57	−0.04	−0.76	−0.83	−0.53
Spring	$\widehat{\beta_h^{Spring}}$	−7.78	−8.04	−8.58	−9.46	−12.72	−12.42	−8.62	−5.70	−3.88	−2.74	−2.33	−2.06
Summer	$\widehat{\beta_h^{Summer}}$	−28.76	−28.06	−27.46	−28.94	−35.60	−36.55	−31.58	−22.22	−15.86	−12.77	−11.08	−9.87
Regression	$\widehat{\alpha}_h$	0.45	0.46	0.48	0.48	0.40	0.40	0.47	0.62	0.72	0.76	0.77	0.78

Table 3 Cross-sectional correlation matrix of the error terms in the stochastic part of the model for the day-ahead market

Hour	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1	1																							
2	0.96	1																						
3	0.85	0.90	1																					
4	0.86	0.94	0.90	1																				
5	0.83	0.91	0.88	0.98	1																			
6	0.81	0.86	0.82	0.92	0.94	1																		
7	0.76	0.79	0.74	0.83	0.86	0.94	1																	
8	0.66	0.68	0.63	0.69	0.71	0.75	0.84	1																
9	0.58	0.59	0.55	0.60	0.61	0.65	0.73	0.92	1															
10	0.59	0.60	0.56	0.61	0.62	0.65	0.73	0.91	0.97	1														
11	0.63	0.64	0.59	0.64	0.65	0.68	0.76	0.91	0.94	0.97	1													
12	0.64	0.65	0.60	0.65	0.66	0.70	0.77	0.89	0.90	0.94	0.98	1												
13	0.66	0.67	0.63	0.68	0.70	0.73	0.80	0.89	0.88	0.91	0.95	0.98	1											
14	0.66	0.68	0.64	0.70	0.71	0.74	0.81	0.88	0.85	0.88	0.93	0.95	0.99	1										
15	0.66	0.67	0.63	0.69	0.71	0.74	0.80	0.87	0.84	0.87	0.90	0.93	0.96	0.98	1									
16	0.63	0.64	0.60	0.66	0.67	0.71	0.76	0.84	0.81	0.83	0.86	0.89	0.92	0.94	0.96	1								
17	0.56	0.58	0.54	0.59	0.60	0.63	0.68	0.79	0.80	0.81	0.82	0.83	0.85	0.84	0.87	0.92	1							
18	0.50	0.51	0.48	0.52	0.53	0.56	0.61	0.75	0.73	0.75	0.76	0.77	0.80	0.79	0.81	0.86	0.93	1						
19	0.53	0.53	0.50	0.54	0.55	0.58	0.65	0.74	0.72	0.75	0.78	0.79	0.81	0.82	0.83	0.87	0.88	0.92	1					
20	0.52	0.53	0.51	0.55	0.57	0.60	0.65	0.68	0.64	0.67	0.72	0.75	0.78	0.80	0.81	0.82	0.80	0.82	0.92	1				
21	0.54	0.55	0.53	0.59	0.61	0.64	0.69	0.67	0.64	0.66	0.71	0.75	0.78	0.80	0.81	0.80	0.74	0.73	0.84	0.92	1			
22	0.52	0.55	0.53	0.60	0.62	0.65	0.70	0.64	0.61	0.62	0.66	0.69	0.73	0.75	0.76	0.75	0.70	0.67	0.75	0.85	0.96	1		
23	0.52	0.56	0.54	0.61	0.63	0.64	0.66	0.58	0.54	0.56	0.59	0.62	0.67	0.69	0.69	0.67	0.63	0.59	0.66	0.74	0.86	0.94	1	
24	0.42	0.47	0.45	0.53	0.55	0.57	0.57	0.49	0.45	0.47	0.50	0.52	0.57	0.58	0.58	0.57	0.53	0.50	0.55	0.62	0.74	0.83	0.92	1

3.2 The balancing market

For the balancing market, we consider premiums rather than prices. In an optimization model for multi-market trade, it is the difference between the prices that are important for the trading strategy, i.e. if there is a premium over the day-ahead prices to be gained by participating in the balancing market. By analyzing the data for prices, it was found that the correlation between the day-ahead price and the premiums in the balancing markets is much weaker than the correlation between prices in these markets. The weak correlation between day-ahead prices and balancing premiums enables us to model day-ahead prices and the premiums by independent models.

Several methods for forecasting the balancing market are tested in [13]. The balancing market is event-driven, where hours with regulation in either direction are followed by hours of no regulation in an irregular pattern. This behavior cannot properly be modelled by a regular time-series model, since the timing between nonzero-observations is unevenly spaced. The balancing market only really exist in the time steps where there is demand in the market, i.e. when there is an unbalance. The size of demand is the amount of power that is needed to bring the system back in balance. The amount needed to bring the system back in balance may be less than the production capacity of individual agents, so there is a limitation on the amount that can be sold in the balancing market.

Due to this limited volume in balancing market, we model the trade limit in this market as well as the premiums. That is, for each time step, the maximum volume that may be traded in the market is limited by an upper bound. This upper bound is a stochastic parameter in the multi-market optimization problem, see Sect. 2. The balancing markets is thus described by two time series, one for premiums and one for the volume or trade limit. We model these together using one of the methods proven to have good performance in [13], namely the method of unevenly spaced time-series based on the work of [17].

The key for unevenly spaced time-series is to distinguish between the probability of the arrival of demand and the size of the demand. We define demand for balancing to exist whenever there is a need for either up or down regulation. In the training period, there is demand for balancing in 71% of the hours. The average time between demand is 1.23 h. We use the approach of [18] and model the time between arrival of demand as a moving average process where π_t is the (moving) average time between arrivals and q_t is the number of time steps since the last event. If we denote the balancing volume in hour t as V_t^{BM}, then

$$\pi_t = \begin{cases} \pi_{t-1} & \text{if } V_t^{BM} = 0 \\ \pi_{t-1} + \alpha(q_{t-1} - \pi_{t-1}) & \text{if } V_t^{BM} \neq 0. \end{cases} \tag{9}$$

The probability of balancing for each hour is then $1/\pi_t$. The parameter α needs to be estimated, and this is done by minimizing the sum of square of residuals from the empirical arrival rate from historical data and the estimated arrival rate. The optimal value was found to be close to 0.01. The model for the arrival rate is not complete without information of the balancing volume, V_t^{BM}, as this volume is needed in each hour to update the moving average process for probability of demand. The balancing

volume is modelled as a stationary unevenly spaced time series of order 1, using the methods in [19]. That is, if Δ is the irregular time between non-zero values, the balancing volume in hour $t + \Delta$ is

$$V_{t+\Delta}^{BM} = \theta^\Delta V_t^{BM} + \sigma_\Delta + \epsilon_{t+\Delta} \tag{10}$$

where $\epsilon_{t+\Delta} \sim N(0, 1)$ and $Cov(\epsilon_t, \epsilon_s) = 0$ for every $t \neq s$ and $\sigma_\Delta^2 = \sigma^2 \frac{(1-\theta^{2\Delta})}{(1-\theta^2)}$ for some $\sigma > 0$. The parameters θ and σ must be estimated, again using the method in [19]. For our data set, we obtain values for $\hat\theta = 0.83$ and $\hat\sigma = 70.54$. Our combined model for state determination and volume is thus an unevenly spaced time series, where the arrival rate is specified by a moving average process.

Notice that with this method, the moving average process will determine, in each time step, if there is demand for regulation or not. However, this does not give us any information about whether the demand is for up or down regulation. The direction of regulation is thus determined by the sign of the balancing volume, which may take both positive and negative values. The above model neither say anything about premiums. However, the balancing market premium is only relevant when there is a demand for balancing power. Thus, the premium for up regulation is only needed in the hours where there is a volume for up regulation, and vice versa for down regulation. In the modelling, we therefore use the same moving average process that determines the arrival of demand for regulation to determine the "demand" for regulation premiums. In this manner, only the hours with up or down regulation will have a corresponding premium. The AR process that determines the size of the premium is however fitted separately. That is, we determine the coefficients $\hat\theta$ and $\hat\sigma$ of the equivalent of Eq. (10) for the balancing premiums. We estimate the parameters using the same methods as for the balancing volume. The obtained values are $\hat\theta = 0.76$, and $\hat\sigma = 38.58$.

The historical data used to fit the time-series for the balancing market describes the total volume of activated power in the NO2 area. We must make assumptions on how much of the total volume that can be supplied by an individual agent or hydropower system. One approach is to assume that the individual producer may take a percentage of the total market volume, e.g., 10%, in every hour where there is demand. This means that a small volume is available to the producer in most hours. Another approach is to assume that in each hour, the total market volume may be activated from a single producer with a given probability, e.g., one in ten times. The probability may be related to the number of agents in the power market. The latter approach is taken in this work.

4 Scenario generation

With the specified models, we proceed to forecasting future values and scenario generation. The method used for scenario generation is described in [9], and is based on combining point forecasts with historical forecast errors, i.e., the difference between historical forecasted values and historical realized values. For this reason, we divide our data into 3 sets: a training set for the time series models, a training set for the

scenario-based forecast method (where we find the forecasting errors), and a final set where we test the scenario generation method, see Fig. 2. All data is taken from Nord Pool's ftp server and include day-ahead prices, balancing market prices and turnovers. In addition, we also make use of historical exchange rates between Euros and NOK. The training set for the time-series models are the years 2014–2016. The next period is the training period for the forecast-based scenario generation method, which is taken to be data for 2017. This period may also be seen as the test period for the forecasting models. For each day in this period, we use the time-series models to forecast prices and volumes in the markets. We then compare the forecasted values with realized values to find historical forecast errors. The historical forecast errors are used as input for generating scenarios for the final data set. In this period, we generate scenarios by computing daily point forecasts and combining this with the historical forecast errors. The test period for scenario generation is the first 33 weeks of 2018.

In this work, we use a 72-h forecast horizon. This corresponds to the time-horizon of the optimization model. Usually, short-term hydropower scheduling have a horizon of about one week ahead. This is to avoid end-of-horizon effects and for consistent linking to longer-term models. The 72-h horizon is used here for computational reasons.

As stated earlier, we use daily branching. Specifically, this means that the scenario trees have branching points at the end of Hour 24 and 48. The tree structure is illustrated in Fig. 1 when the tree branches into 2 new scenarios at each branching step. In this section, we make individual scenario trees for each market, see Fig. 3. For the balancing market, positive values represent up regulation while negative values represents down regulation (Fig. 4).

In stochastic programming models, the variation or distribution of the uncertain parameters is important. It is the variation of prices we want to describe in the scenario tree. Therefore, we would like to test the probabilistic quality of the generated scenarios. For this we use a percentile-based test similar to the approach used in [9]. That is,
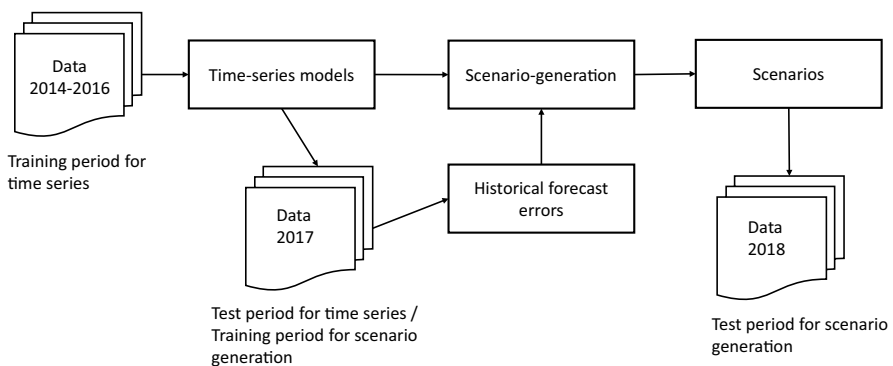
Fig. 2 An illustration of how the data is divided into test and training sets for the different steps of the modelling process
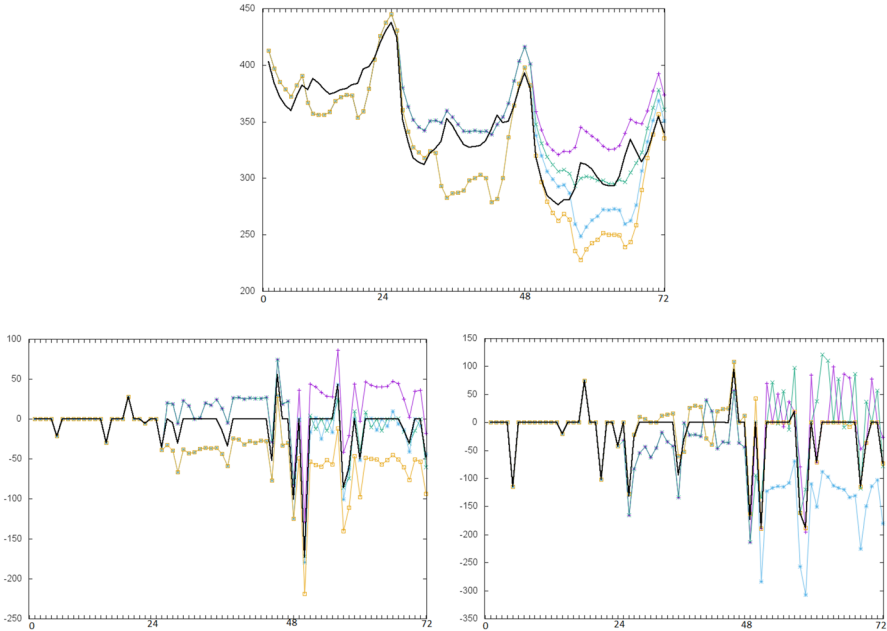
Fig. 3 Examples of the generated scenario values for day-ahead prices (**a**), the balancing market premium (**b**) and the balancing market volumes (**c**). The black line is the point forecast

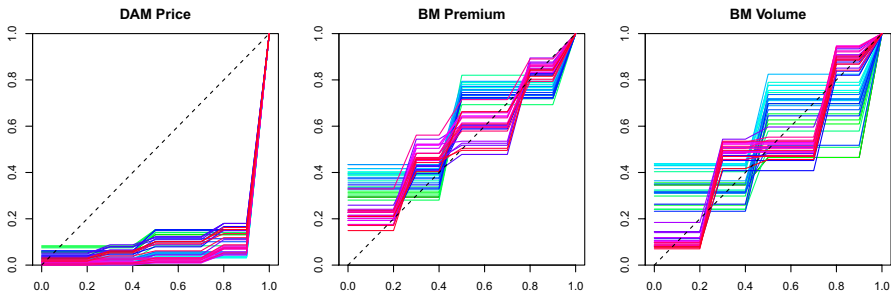**Fig. 4** Probability plots showing actual vs. expected percentiles of the generated scenarios when the tree has 2*2=4 scenarios. Each line represents one forecast length t

for forecast length Δt on each day d in the test period for scenario generation, we compute the percentile of the observed value within the generated scenarios:

$$p_{d\Delta t} = |\{s \in 1 \dots S_{\Delta t} : x_t^s \leq x_{\Delta t}\}| / |S_{t\Delta t}| \tag{11}$$

where $x_{\Delta t}^s$ are the scenarios values and $x_{\Delta t}$ are the actual observed values. This means that $p_{d\Delta t}$ is zero/one if the value is lower/higher than in the scenarios. If the scenarios represented the uncertainty perfectly, we should see $p_{d\Delta t} \leq q$ in $100q$ percent of the cases, for each Δt. To test this, we compute the frequencies of the quantiles,

$$F_{\Delta t} = |\{d \in D : p_{d\Delta t} \leq q\}|/|D|, \tag{12}$$

for each q and plot this against q in a probability plot, see Fig. 2. In the plots, the horizontal axis is q, the vertical axis is $F_{\Delta t}$, and each line is $F_{\Delta t}$ for individual forecast lengths Δt. If the percentiles of the scenarios perfectly match the observed values, the lines should be close to the optimal line between $(0, 0)$ and $(1,1)$.

For day-ahead prices, the percentiles of the scenario values deviate quite a bit from the optimal line. Note that, since the scenario generation method adds noise to an externally provided point forecast, the quality of the scenarios depends on the quality of the time-series model as well as the scenario generation method. The QQ-plot for the day-ahead prices indicates that the model for day-ahead prices is not very good because we are constantly underpredicting the true percentiles of the distribution. The reason for this is that the day-ahead prices were higher than normal in 2018, see Fig. 5 that shows the realized prices and the first-day forecast for the day-ahead prices in the first 33 weeks of 2018. From the plot, it is clear that the model underpredicts the realized values. Even though the day-ahead model shows poor performance, we use it to generate input for the tests presented in Sect. 5. The objective here is not to make an advanced forecasting model for the day-ahead market, but rather to test the multi-market bidding model. The tests are valid as long as all the tests use the same (poor) model for the day-ahead market.
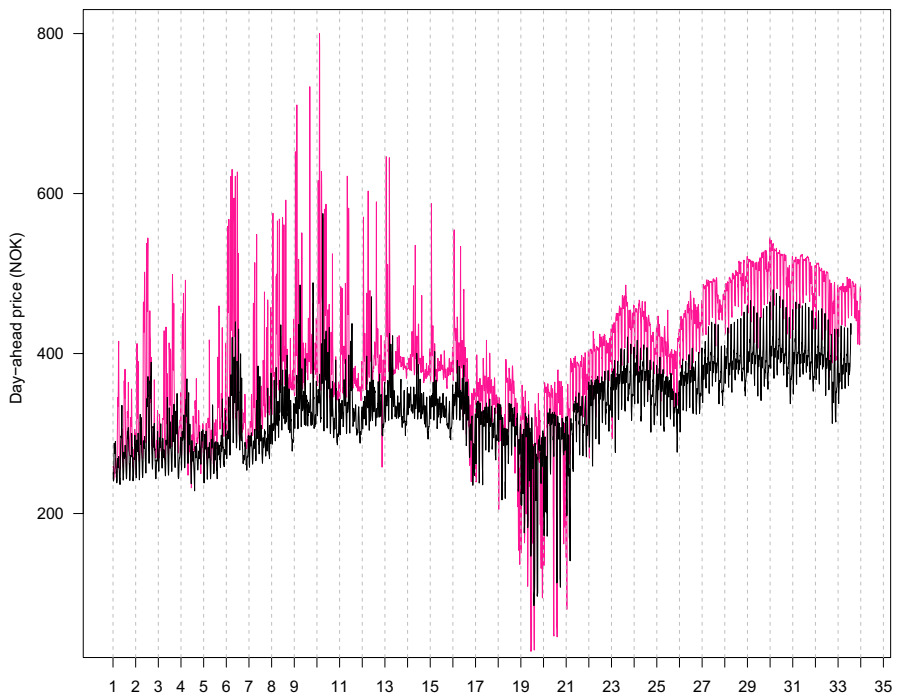
Fig. 5 The realized prices (pink) and first-day forecast (black) for the day-ahead prices in 2018

Next, observe that there are several intervals of constant values on the curves in Fig 2. These occur because the scenarios trees used here only have 4 scenarios in total. In fact, for the shortest time horizons (within the second day), the scenarios only take on two possible values. This corresponds to the blue and green lines in Fig. 2. Since the scenarios are equiprobable, the values of $p_{d\Delta t}$ must be multiples of 0.5 here. For the third day, there are 4 possible scenario values and the red and pink lines have steps that are multiples of 0.25.

The stepwise curves in Fig. 2 illustrates that 4 scenarios (i.e. branching into 2 new scenarios at each branching step) only provides a rough approximation to the true uncertainty of prices and volumes in the markets. Using more scenarios would allow the approximation to become closer to the optimal diagonal line as the steps would be smaller, see Fig. 6 where the scenario trees have 9 scenarios in total . However, for large watercourses consisting of several reservoirs and power plants, the size and complexity of the optimization model is large, even in the deterministic case. Including multiple products and their stochastic price and volume adds new layers of complexity. In addition, in the above presentation, we have generated separate scenario trees for the day-ahead price and the balancing market price and volume. Before being used in the optimization model, these separate scenarios trees are combined all-against-all to generate a total tree. If the scenario tree for the day-ahead prices has 4 scenarios and the scenario tree for the balancing market has 4 scenarios, the total tree will have 16 scenarios. If each of the markets have 3 scenarios, the total tree will have 81 scenarios. The size of the scenario tree will be a trade-off between accurate representation of the markets and computational time.

5 Case study

We illustrate use of the multi-market model by analyzing the value of participating in the balancing market and the value of coordination across markets. We run several cases to evaluate these gains:

1. **Day-ahead only:** Optimize the trades in the day-ahead market in an optimization model that only sees the day-ahead market.
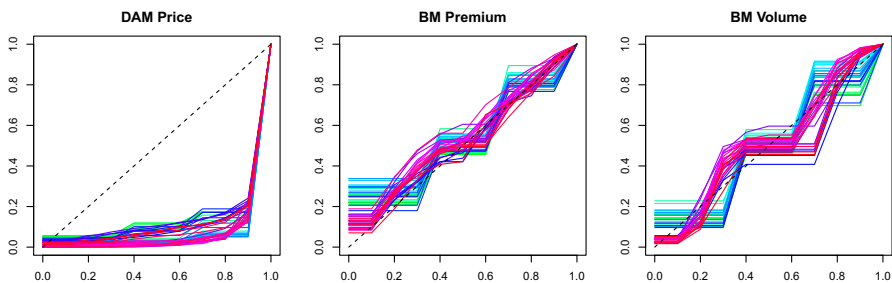

Fig. 6 Probability plots showing actual vs. expected percentiles of the generated scenarios when the tree has 3*3=9 scenarios. Each line represents one forecast length t

2. **Sequential bidding:** Optimize trades in the day-ahead market as if this was the only market (same as in case 1). Lock the decisions for the day-ahead market and then determine the balancing market trades in an optimization that sees both the day-ahead and balancing market.
3. **Coordinated bidding:** Jointly optimize trades in the day-ahead and balancing markets in an optimization model that sees both markets.

The above three cases are first run with the market non-anticipativity constraints in Eq. (5) applied to both the day-ahead market trades and the balancing market trades. This means that the balancing volumes must be determined at the same time as the day-ahead trades, i.e. with limited information. Notice that this does not necessarily mean trades are coordinated across markets. In the sequential approach, the day-ahead market trades are locked to the values obtained when optimizing without seeing the balancing market. The balancing market volumes may then be optimized against uncertain prices and volumes. The effect of including Eq. (5) for the balancing trades is namely that the balancing market trades must be determined before knowing the balancing market prices and volumes. The cases including Eq. (5) will give a lower bound on the value of participating in the balancing market, because the balancing market trades have to be determined the day before instead of 45 min before delivery.

If, however, the constraints in Eq. (5) are applied only to the day-ahead trades, the trades in the balancing market may be determined in real-time, i.e. with perfect information about the prices and volumes in the balancing market. Again, this does not necessarily mean that the trades are coordinated. In the sequential approach, the day-ahead trades cannot be changed to accommodate favorable opportunities in the balancing market, even if the balancing trades can be optimized against known prices and volumes. The cases not including Eq. (5) for the balancing trades will thus give an upper bound on the value of participating in the balancing market because the balancing market trades may be determined with perfect information.

In total, 6 different runs of the model are executed each test day; the 3 cases in the list above for limited and perfect information about the balancing market. Tests are performed for 46 instances corresponding to market input for different indivudial days in the first 33 weeks of 2018.

The analysis is done for an example hydropower system, where a plant with two generating units is located downstream of a single reservoir. The storage capacity of the reservoir is 50 Mm^3. The maximum discharge is 150 m^3/s, which means that the reservoir may be emptied in about 90 h. This is a very flexible production system, since there is large reservoir capacity and no complicated waterways or other operating constraints. The generators have a minimum production level of 30 and a maximum level of 50 MW.

On each test day, separate scenario trees for the day-ahead market prices and the balancing premiums and volumes are generated using the time-series models and the forecast-based scenario generation method. As a base case, we use 3 scenarios at each branching point (e.g. at the end of Hour 24 and Hour 48), resulting in 9 scenarios for each market. Combining these into a total tree, the optimization model sees

81 scenarios. Each run of the model take about 15 seconds on an Intel 3.00 GHz lab top computer with 64-bit RAM.

The value of storing water in the reservoir is set to the average of the day-ahead prices in the scenario tree for the current day. This means that the optimization problem finds the balance of selling and producing within the optimization period versus storing the water for use in the future. Traditionally, the value of water is determined from longer-term optimization models that consider management of the reservoir over months or years. In hydropower scheduling, the long-term effect of the trading strategy is important because of the inventory problem of storing water over time. Use of the multi-market model should therefore be tested over longer time horizons, and ideally over several years to cover seasonal variations and dry/wet years. However, such aspects are postponed for future work and we use the more simple way of calculating the value of water as the average day-ahead price each day. Also, the initial reservoir level is re-set at 50% filling each day and the tests consider separate days as individual problem instances.

5.1 Limited information

The results for the limited information case is shown in Fig. 7. The figure shows bar plots of the objective function value for each day. The objective function value is made up of the end value (EV) which is the value of water stored in the reservoir at the end of the optimization horizon, income from the day-ahead market (DAM), income from the balancing market (BM) and start-up costs (Start). In the first case when the producer only participates in the day-ahead market, there is no income from the balancing market. The optimization finds a balance between producing within the optimization period and saving water in the reservoir. This balance is
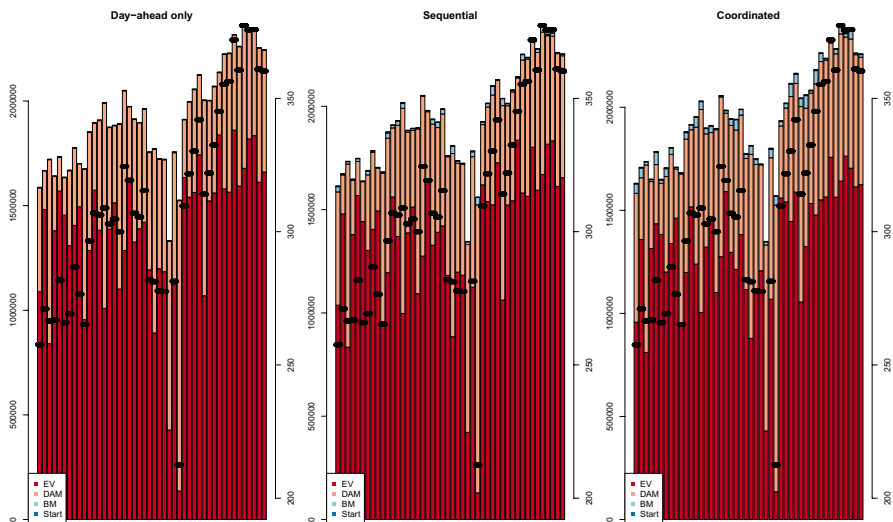

Fig. 7 Bar plots showing the objective function value for each test day. The left y-axis shows the objective function value in NOK. The right y-axis and the black dots gives the water value for each day

determined by the water values (shown as black dots) and the prices in the market. For the two cases involving the balancing market, we see that both the sequential and coordinated approach gives added profits over participating in the day-ahead market only.

To better discern the differences between the 3 cases, consider Fig. 8 which shows the results for 6 selected days. For the cases involving the balancing market, the total objective function is slightly larger and the contribution from each element (EV, DAM, BM and Start) is also different. The effects are small, but in general, when the balancing market is included, more of the objective function value comes from trades. This can be seen from the fact that the balance of producing vs storing water (e.g the size of BM+DAM vs the size of EV) is shifted more towards producing when the balancing market is included. This effect is more pronounced in the coordinated case than in the sequential case. Figure 9 gives boxplots of the percentage gain in objective function value when participating in the balancing market using the sequential and coordinated approach. The rightmost plot shows the added gain of coordination. The average gain of using the sequential approach is 0.8%, while the average gain of the coordinated approach is 1.4%. The extra gain of the coordinated approach over the sequential is on average 1.7%.

5.2 Perfect information

Moving to the perfect information case, Fig. 10 shows the equivalent of Fig. 8. Again, in general, more of the value comes from trades when the balancing market is included. This effect is largest when using the coordinated approach. This is natural because the producer has more freedom in the joint optimization, and, in addition, the balancing market prices and volumes are known because of the
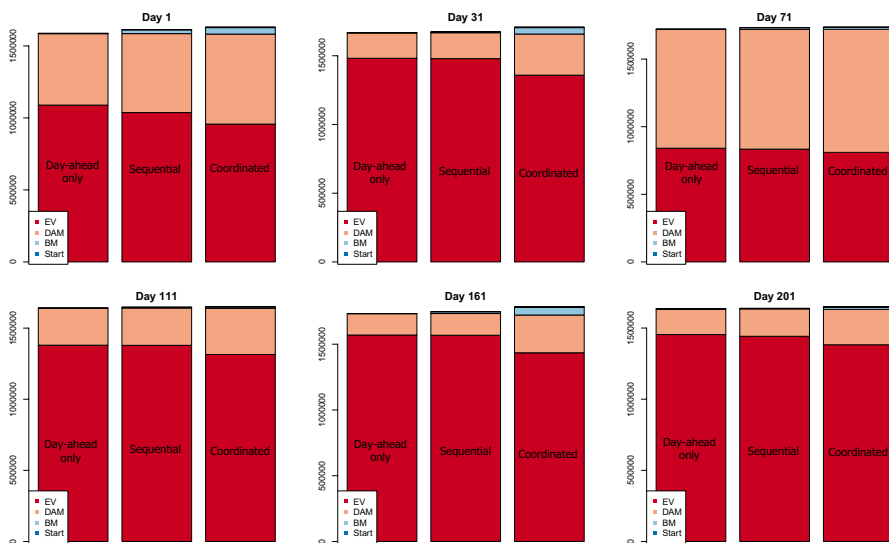

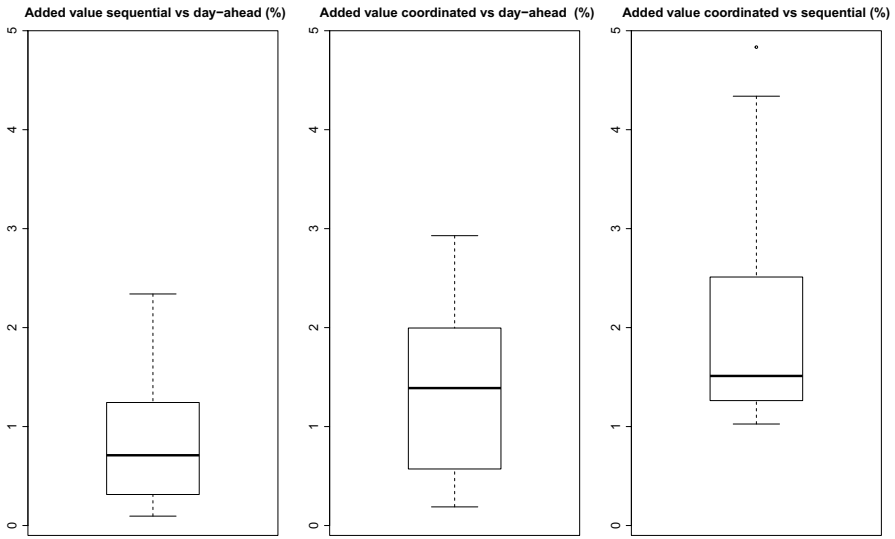Fig. 8 Comparison between the 3 strategies on 6 selected days. Limited information

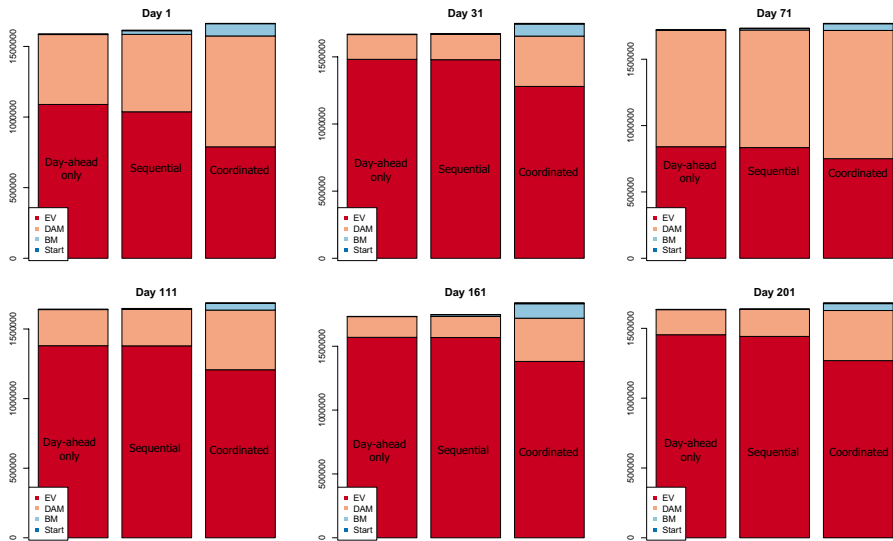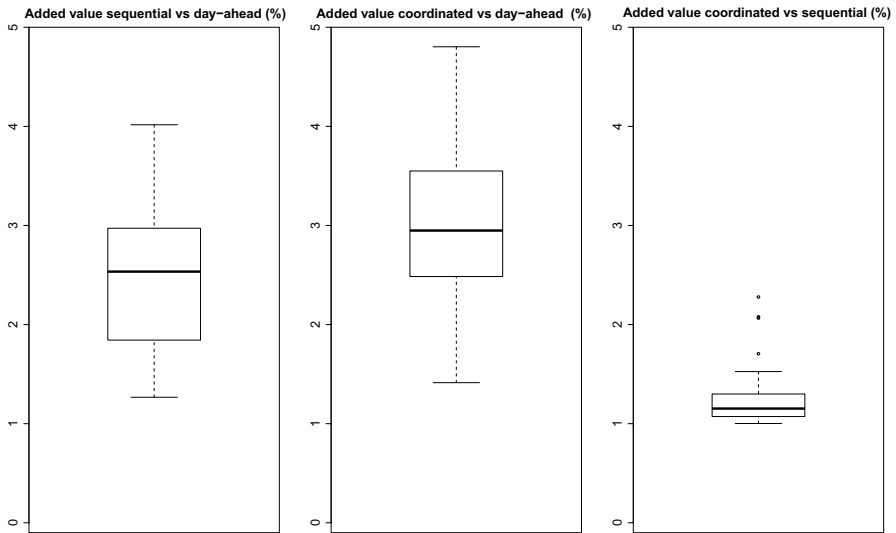**Fig. 9** Box plots of the gains from using different strategies. Limited information

Fig. 10 Comparison between the 3 strategies on 6 selected days. Perfect information

assumption of pefect information. Figure 11 shows boxplots of the gains of using different approaches when having perfect information about the balancing market. The average gain of using the sequential approach is 2.6%, while the average gain of the coordinated approach is 2.9%. The extra gain of the coordinated approach over the sequential is on average 1.1%.

Fig. 11 Box plots of the gains from using different strategies. Perfect information

Summarizing the results for the limited and perfect information case, the value of participating in the balancing market using the sequential approach is between 0.8 and 2.6%. Using the coordinated approach, the producer may gain between 1.4 and 2.9%. So participating in the balancing market using either method provides increased revenues compared to only participating the day-ahead market. The value of coordination, i.e. the value of using the coordinated over the sequential approach, is found to be higher in the limited information case (1.7%) than in the perfect information case (1.1%). This indicates that, the more uncertainty, the higher is the value of coordination.

Even though we calculate lower and upper bounds on the value of participating in the balancing market, the results obtained are only valid for the example hydro-power system being tested. This system is very flexible and the gain might thus be higher than for other river chains with more constrained production systems. Results will be case-specific, but repeating the analysis outlined above will give lower and upper bounds in each case.

Another interesting result is that the balance of producing versus saving water is shifted towards producing in the cases where the balancing market is included. When participating in the balancing market, the producer is able to get a higher price for the water produced. When this is compared to the value of water (which in this case is determined based on the day-ahead prices only) a higher price means that more water should be produced within the optimization period instead of saved for later production. Multiple short-term markets mean that producers may trade their way to the higher prices for their production. The results here indicates that opportunities in multiple markets should be reflected in longer-term models that calculate the value of water.

Table 4 Stability of calculated lower bounds

Day-ahead scenarios	Balancing scenarios	Total scenarios	Sequential vs. day-ahead	Coordinated vs. day-ahead	Coordinated vs. sequential
$2 * 2 = 4$	$2 * 2 = 4$	$4 * 4 = 16$	0.8	1.5	1.9
$3 * 3 = 9$	$3 * 3 = 9$	$9 * 9 = 81$	0.8	1.4	1.7
$4 * 4 = 16$	$4 * 4 = 16$	$16 * 16 = 256$	0.7	1.3	1.9

Numbers in%

Table 5 Stability of calculated upper bounds

Day-ahead scenarios	Balancing scenarios	Total scenarios	Sequential vs. day-ahead	Coordinated vs. day-ahead	Coordinated vs. sequential
$2 * 2 = 4$	$2 * 2 = 4$	$4 * 4 = 16$	2.1	2.9	1.4
$3 * 3 = 9$	$3 * 3 = 9$	$9 * 9 = 81$	2.6	2.9	1.1
$4 * 4 = 16$	$4 * 4 = 16$	$16 * 16 = 256$	2.4	1.8	1.3

Numbers in%

5.3 Stability of the solution

As a final note, the stability of the bounds on the gain of participating in the balancing market is tested with respect to the number of scenarios used. To test for stability, we generate trees of different sizes each representing the same day. The cases described above are repeated for each test day using 2 new scenarios at each branching point and using 4 new scenarios at each branching point. This means that the separate trees for each market will have $2 * 2 = 4$ or $4 * 4 = 16$ scenarios, resulting in 16 or 256 scenarios in the total tree. The cases with 16 scenarios takes only seconds to solve, but using 256 scenarios takes over 1 h for each run of the model. Larger instances are therefore not tested.

The results are summarized in Tables 4 and 5 and shows that the calculated bounds vary for different tree sizes. The general results are however the same: there is a gain from participating in the balancing market using either the sequential or the coordinated approach, and the gain of coordination is smaller when the producer has perfect information.

6 Conclusions

In this paper, we formulate the multi-market bidding problem for power producers as a stochastic program. The formulation is implemented within the framework of optimization models that is used by the Nordic hydropower industry. We also present how input to the stochastic program may be generated by using a forecast-based

scenario generation method combined with time-series models that predicts future prices and turnovers in the markets. The model is applied in a case study to investigate the value of participating in the Nordic day-ahead and balancing market. A producer may participate in the balancing market either by considering trades in the markets sequentially or coordinated.

Using cases with limited and perfect information about the balancing market to calculate lower and upper bounds, we find that the value of participating in the balancing market using the sequential approach is between 0.8 and 2.6%. Using the coordinated approach, the producer may gain between 1.4 and 2.9%. These results show that participating in the balancing market using either method provides increased revenues compared to only participating in the day-ahead market. The value of coordination, i.e. the value of using the coordinated over the sequential approach, is found to be higher in the limited information case (1.7%) than in the perfect information case (1.1%). This indicates that, the more uncertainty, the higher is the value of coordination. The results obtained are case-specific and only valid for the example hydropower system being tested. The method of using limited and perfect information to calculate bounds may however be applied to other systems.

Further work on the topic of coordinated bidding should include efforts to improve the modelling of prices and volumes in the various markets, as well as dependencies and correlations between them. To test the long-term effect of using different trading strategies or forecasting methods, simulation or backtesting frameworks could be developed to evaluate and compare the performance of different modelling choices and assumptions. In addition, investigation could be made into also including the stochastic nature of activation of reserves or balancing power in the formulation of the optimization model.

References

1. SINTEF Energy Research SHOP website. https://www.sintef.no/en/software/shop. Online; Accessed 8 Apr 2018
2. Boomsma, T.K., Juul, N., Fleten, S.-E.: Bidding in sequential markets: The Nordic case. Eur. J. Oper. Res. **238**(3), 797–809 (2014)
3. Birge, J.R., Louveaux, F.: Introduction to Stochastic Programming. Springer, New York (2011)
4. Fleten, S.-E., Pettersen, E.: Constructing bidding curves for a price-taking retailer in the norwegian electricity market. IEEE Trans. Power Syst. **20**(2), 701–708 (2005)

5. Ugedo, A., Lobato, E., Franco, A., Rouco, L., Fernàndez-Caro, J., Chofre, J.: Strategic bidding in sequential electricity markets. IEE Proc. Gen. Transm. Distrib. **153**, 431–442 (2006)
6. Faria, E., Fleten, S.-E.: Day-ahead market bidding for a Nordic hydropower producer: taking the Elbas market into account. CMS **8**(1), 75–101 (2011)
7. Klæboe, G.: Stochastic short-term bidding optimisation for hydro power producers, Doctoral Thesis, NTNU, Norway (2015)
8. Vardanyan, Y.: Optimal bidding of a hydropower producer insequential power markets with risk-assessment: Stochastic programming approach. Doctoral Thesis, KTH, Sweden (2016)
9. Kaut, M.: Forecast-based scenario-tree generation method. Optimization Online (2017)
10. Heitsch, H., Römisch, W.: Scenario tree modeling for multistage stochastic programs. Math. Program. **118**(2), 371–406 (2009)
11. Skjelbred, H.I.: Unit-based Short-term Hydro Scheduling in Competitive Electricity Markets, Doctoral Thesis at Norwegian University of Science and Technology (2019)
12. NordPool annual report 2017. https://www.nordpoolgroup.com/globalassets/download-center/annual-report/annual-report-nord-pool_2017.pdf. Online; Accessed 4 Sep 2018
13. Klæboe, G., Eriksrud, A.L., Fleten, S.-E.: Benchmarking time series based forecasting models for electricity balancing market prices. Energy Syst. **6**(1), 43–61 (2015). https://doi.org/10.1007/s12667-013-0103-3
14. Huismann, R., Huurmann, C., Mahieu, R.: Hourly electricity prices in day-ahead markets. Energy Econ. **29**, 916–928 (2007). https://doi.org/10.1016/j.eneco.2006.08.005
15. Klæboe, G.: Forecasting hourly electricity prices for bidding optimization. In: 12th International Conference on the European Energy Market (EEM), (2015). https://doi.org/10.1109/EEM.2015.7216636
16. https://cran.r--project.org/web/packages/systemfit/index.html
17. Croston, J.D.: Forecasting and stock control for intermittent demands. Oper. Res. Q **23**, 289–303 (1972). https://doi.org/10.2307/3007885
18. Willemain, T.R., Smart, C.N., Shockor, J.H., DeSautels, P.A.: Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. Int. J. Forecast. **19**, 529–538 (1994). https://doi.org/10.1016/0169-2070(94)90021-3
19. Erdogan, E., Ma, S., Beygelzimer, A., Rish, I.: Statistical models for unequally spaced time series. In: Proceedings of the 2005 SIAM International Conference on Data Mining, (2005). https://doi.org/10.1137/1.9781611972757.74