SINTEF

TR A7559- Unrestricted

Report

Hydropower Scheduling Considering Energy and Reserve Capacity Markets

Using a combined SDP/SDDP algorithm

Author(s)
Arild Helseth
Birger Mo
Marte Fodstad

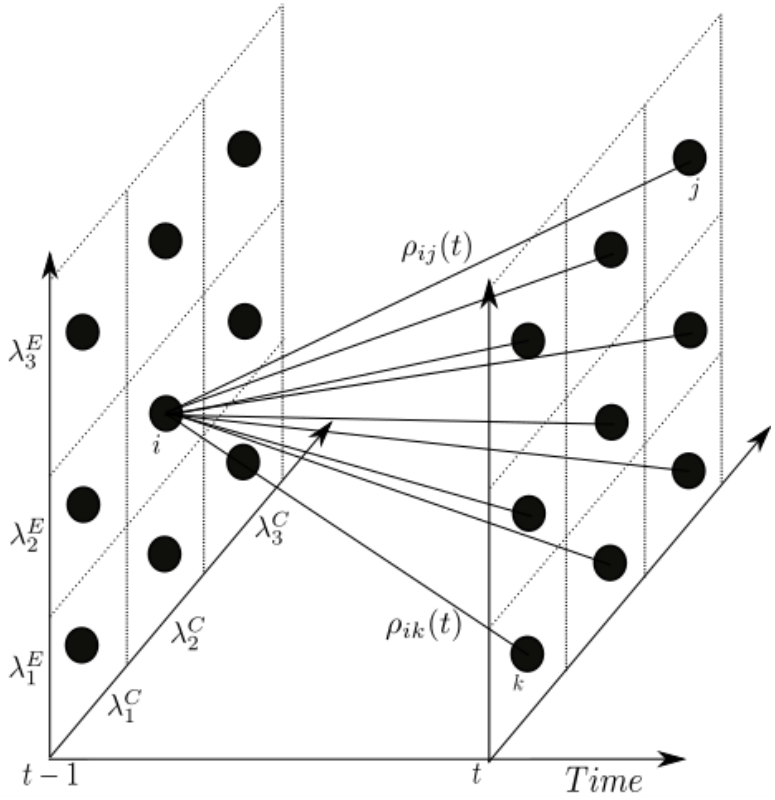

SINTEF Energy Research
Production Planning
2016-05-23

SINTEF

SINTEF Energi AS
SINTEF Energy Research

Address:
Postboks 4761 Sluppen
NO-7465 Trondheim
NORWAY

Telephone:+47 73597200
Telefax:+47 73597250

energy.research@sintef.no
www.sintef.no/energi
Enterprise /VAT No:
NO 939 350 675 MVA

KEYWORDS:
Hydropower scheduling
Reserve capacity
Stochastic optimization
Linear programming

Report

Hydropower Scheduling Considering Energy and Reserve Capacity Markets

Using a combined SDP/SDDP algorithm

VERSION
3

DATE
2016-05-23

AUTHOR(S)
Arild Helseth
Birger Mo
Marte Fodstad

CLIENT(S)
Client(s)

CLIENT'S REF.

PROJECT NO.
502000395

NUMBER OF PAGES/APPENDICES:
41 + Appendices

ABSTRACT

We present a method for medium-term hydropower scheduling considering both sales to an energy and reserve capacity market. The method is implemented in a prototype model and is founded on the same combined SDP/SDDP algorithm that is used in ProdRisk.

It is elaborated both analytically and numerically how consideration of the reserve capacity markets impacts the computed water values. We also discuss how the model results can be used to price reserve capacity.

PREPARED BY
Arild Helseth

SIGNATURE
Arild Helseth

CHECKED BY
Michael Belsnes

SIGNATURE
Michael Belsnes

APPROVED BY
Knut Samdal

SIGNATURE
Knut Samdal

REPORT NO.	ISBN	CLASSIFICATION	CLASSIFICATION THIS PAGE
TR A7559	978-82-594-3655-9	Unrestricted	Unrestricted

Document history

VERSION	DATE	VERSION DESCRIPTION
1	2016-03-17	Submitted to steering group for discussion.
2	2016-04-11	Revised version for internal quality control
3	2016-05-23	Final version

SINTEF

Table of contents

A Nomenclature

1 Introduction

Today's operational models for long- and medium term hydropower scheduling consider sales of energy as the only opportunity for the producer to earn money. However, as the producers need to adapt to a future with increasing share of renewable intermittent generation and emphasis on harmonising various power markets on European level, so do the scheduling models.

In the research project "Integrating Balancing Markets in Hydropower Scheduling Methods" the key research question is to find how the valuation of water is affected when considering balancing markets. The project broadly defines balancing markets to comprise markets for intra-day, reserve capacity and regulating power. In theory, all relevant markets should be included when computing the optimal strategy for a hydropower system, but that would be an exhaustive computational task. Keep in mind that the current scheduling models, e.g., ProdRisk and Vansimtap, involve uncertainty in inflow and electricity (energy[1]) prices and apply weekly decision stages over a period of analyses of typically 3-5 years.

The balancing markets generally allow trading two different products: energy and reserve capacity. It is questionable how the incorporation of additional markets cleared at different time scales involving the same product (energy) would contribute to the strategic scheduling of reservoirs (represented by water values). Thoroughly assessing that impact would call for a model with fine time resolution and many decision stages within the week. In our opinion, incorporating the possibility of selling reserve capacity would potentially have a larger impact on the water values. The work presented in this report therefore focuses on the two different products energy and reserve capacity in a generic manner, without going into too much details about the different markets in which each of those products are traded.

In the Nordic market, reserve capacity is procured by the TSO's as several different products. Normally, one separates between primary (FCR), secondary (FRR-A) and tertiary (FRR-M) reserves. Today, each country's TSO primarily buys the reserve capacity products separately, but there is a political aim to further coordinate the provision of reserve capacity. In most European market designs, reserve capacity is primarily procured before clearing the day-ahead market. In the Norwegian case, Statnett buys all three reserve-types the week ahead of operation[2]. This is done to ensure that sufficient reserve capacity is available and is not bid into the energy market(s). Please see [5] for more information on the Nordic power market designs.

This report presents a method for including sales of capacity into an existing hydropower scheduling algorithm. The existing algorithm is a combination of stochastic dynamic programming (SDP) and stochastic dual dynamic programming (SDDP), and is similar to what is used in ProdRisk. In the presented model, we treat reserve capacity as generic product, and the method is therefore flexible regarding the type of reserve capacity being traded. We assume that the producer is a risk-neutral price-taker in both the energy and reserve capacity market. In the analyses we have assumed that the capacity should be spinning (rotating) and symmetric (same amount for up- and down-regulation). In line with the current market design, we require the sales of capacity to take place before knowing the energy price.

[1] We use the term *energy price* (rather than electricity price) throughout this document in order to clearly separate the two products energy and capacity.
[2] With the exception of FCR, which is also bought after clearing the day-ahead market.

2 Model description

In the following we describe the basics of the SDP/SDDP algorithm, similar to that used in ProdRisk. Furthermore, we describe in detail how sales of reserve capacity can be included in this algorithm. The resulting mathematical model was presented in compact form in [6]. A preliminary version was presented in [7]. For further reading about the algorithm applied in ProdRisk, see [4, 2, 3].

The mathematical model was implemented in the C++ programming language, and requires access to a third-party linear programming (LP) solver such as COIN Clp, CPLEX or Gurobi.

2.1 Problem solving by decomposition and sampling

The objective of the medium-term hydropower scheduling problem is to find the release policies that maximize the expected profit over the period of analyses. It is normally solved for a period of 2-5 years starting at a known initial state[3]. It is important to include uncertainty both in future inflow to the reservoirs and electricity prices.

We let the model take decisions in weekly stages. That is, for each week the value of the uncertain variables for that week are known to the model, and the optimal release policy can be made for that week.

The structure of decision sequences can be represented in a scenario tree. For each decision node the tree further branches into a set of new decision nodes in each decision stage. Let's say we plan for a period of 104 weeks and have 12 branches each week, the full scenario tree would have 1.72×10^{112} nodes, which is far beyond what can be stored in memory and be solved in reasonable computation times. For this reason we rely on algorithms that:

a) Decompose the optimization problem into weekly decision problems that can be solved independently;
b) Apply sampling algorithms to avoid searching the entire tree, and;
c) Apply cut sharing to efficiently "collapse" the scenario tree.

A main iteration of the SDDP algorithm is illustrated in Figure 1. It consists of a forward and backward iteration, as briefly outlined below. Note that this illustration and the explanation below refer to the SDDP-part of the combined SDP/SDDP model used e.g. in ProdRisk. Stochastic price is treated by an outer SDP-loop, will be described in section 2.2.

In the *forward iteration* of the SDDP algorithm, we sample a set of inflow scenarios {s1-s3} and simulate week-by-week along these scenarios by solving an LP problem for each week. All scenarios share the initial state, which is defined by a vector of initial reservoirs and inflows in the first week. The simulated state trajectory along the sampled scenarios is shown as the thick black line in Figure 1. We keep track of the simulated reservoir levels and find the expected simulated profit over these scenarios, and that value serves as a lower bound.

In the *backward iteration* cuts are built for each time stage, and added to a list of cuts representing that stage. Consider the state obtained in scenario 1 in stage T-1 in Figure 1. The simulated state trajectory (reservoir level and inflow) up to this point is known from the forward iteration. We now sample 3 vectors of errors or "white noise", so that inflow samples for the coming week can be calculated. For each sample, we solve an

[3] Often referred to as "parallellsimulering"

LP problem. The average dual values on state variables are used to create a cut constraining the expected future profit seen from previous decision stage (T-2). An important feature here is cut sharing; Cuts computed for a specific initial state at a given stage can be shared among all states for that stage.

We use an autoregressive inflow model with lag-1 both in the forward and backward iterations. That means that the inflow for the current week can be found partly by the inflow from the previous weeks, and partly by the sampled white noise. Note that the forward iteration is different in ProdRisk, where historical inflows are used directly.

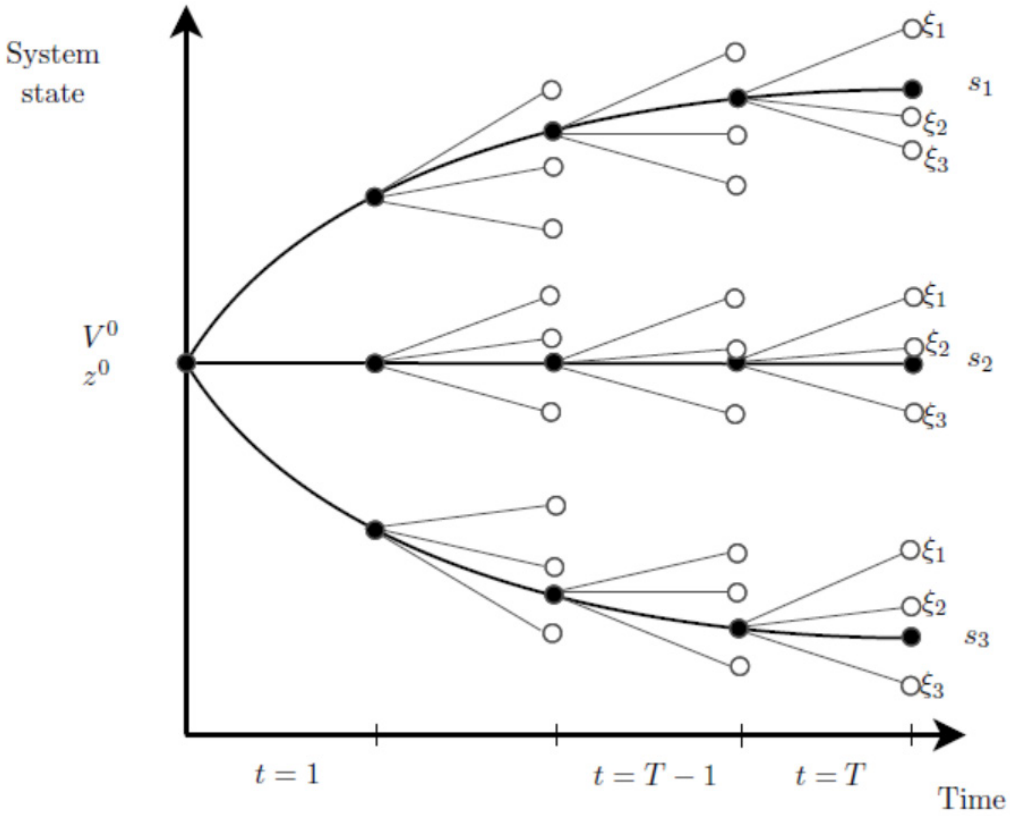

Figure 1 Illustration of a main iteration in the SDDP algorithm.

2.2 Treating stochastic exogenous prices

So far we have focused on the SDDP-part of the algorithm and the treatment of inflow as a stochastic variable. A stochastic exogenous price cannot be treated in the SDDP algorithm as was done with the inflow. This would violate the convexity requirement of the algorithm. For this reason, we will here treat the price process using traditional stochastic dynamic programming and combine it with the SDDP algorithm. This combined SDP/SDDP algorithm was presented in [4, 2], and is used in ProdRisk.

The price process is modelled as a Markov chain using discrete states (price nodes), as illustrated in Figure 2. In each decision stage, cuts are stored per price node. Standing in price node *i* in stage *t-1* in Figure 2, all possible transitions to price nodes in stage *t* are considered when creating the cut.

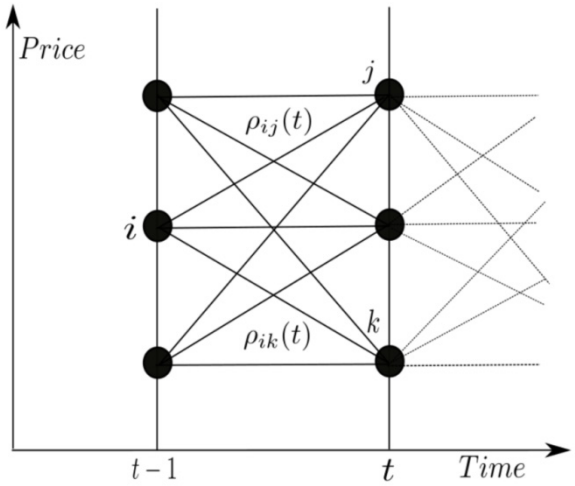

Figure 2 Treatment of the price process in the combined SDP/SDDP scheme.

The price node values and transition probabilities can be generated from a set of price scenarios. Such scenarios are typically obtained from a fundamental market model, such as the EMPS model [8].

The discrete price process illustrated in Figure 2 is embedded in the continuous SDDP scheme illustrated in Figure 1. In the forward iteration the set of sampled scenarios will now include both price-node trajectories and sampled inflow scenarios. In the backward iteration cuts are built and stored for each time stage and price node. That is, cuts are not shared among price nodes within the same time stage. For a thorough mathematical description of this combined SDP/SDDP algorithm, please refer to [2]. Again, we point out that ProdRisk, when used in standard mode, differs from the prototype model presented here in that both prices and inflows takes "observed" rather than sampled values in the forward simulation.

2.3 Time resolution and decision sequences when including reserve capacity

The previous sections gave a short introduction to the overall decomposition method and the treatment of stochastic variables. Now we will focus on the formulation of the (decomposed) weekly decision problem, and see how sales of reserve capacity can be modeled. Before going to the mathematical formulation, the time resolution and decision sequences related to the decomposed problem is discussed.

The decomposed problem represents a weekly decision problem where realizations of stochastic variables are known. That is, when solving the decision problem for a week t, we know the average power price and the accumulated inflow for that week. In order to respect the operational constraints in the watercourse, the week is divided into time steps.

Figure 3 illustrates the time resolution and decision sequences that we have used in this model. A week t is divided into K time steps. In the figure we have divided each day in to three steps, so that K=21.

Energy is sold per time step at a unique price. This price is the product of the average power price per week (represented by a price node) and a pre-defined deterministic relative price profile within the week[4]. Capacity is sold per block, where one block typically covers a set of time steps for weekdays and/or weekend. As an example a block can cover hours 0-8 on weekdays and will then relate to 5 time steps (k=1, k=4, and so on).

Reserve capacity $c_{b,t}$ for block b in week t is sold in the previous week t-1, as indicated in Figure 3. This capacity is sold per watercourse assuming that we know the probability distribution of the reserve capacity prices for week t.

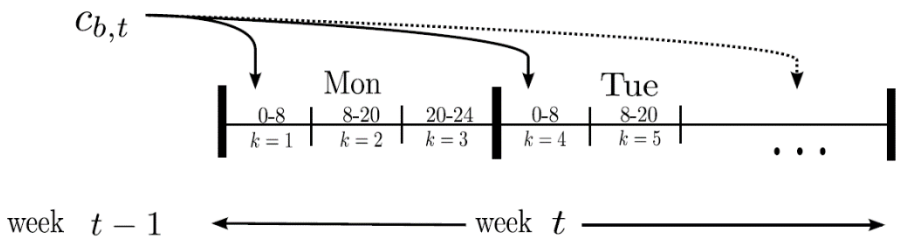

Figure 3 Illustration of time steps within the week. Capacity is sold the week ahead (t-1) directed at a specific block b.

2.4 Treating reserve capacity price and volumes

Markets for reserve capacity have a short history and are continuously subject to changes. Moreover, there is only one buyer (the TSO) and a limited need for capacity. Thus, it is challenging to predict both prices and market volumes for the future in such markets. These quantities will also depend on the type of reserves being traded. While primary and secondary reserves are typically priced higher in the low-load season due to the requirement for sufficient amount (and geographical spread) of spinning capacity, tertiary reserves are only traded in the winter season.

In principle, we would suggest using a fundamental market model to provide forecasts of both energy and reserve capacity prices. The challenge is that the reserve capacity markets have so far not been represented in fundamental models, such as the EMPS model. In the article [6] we used EMPS and formulated a reserve requirement for power producing units in a specified sub-area, and interpreted the dual value of this requirement as the reserve capacity price. We then obtained energy and reserve capacity price series corresponding to the same inflow years simulated in the EMPS model. Note that this approach is likely to underestimate the actual costs of providing reserves, since many relevant constraints are omitted or approximated in the EMPS model.

Once the reserve capacity price is available, one can extend the discrete Markov chain presented in section 2.2 with the reserve capacity price dimension to represent uncertainty in the reserve price. This is illustrated in Figure 4, and further discussed in [6].

[4] This deterministic price profile within the week is found as the average profile considering all price scenarios in ProdRisk.

We assume that the producer is a price-taker in both the energy and reserve capacity markets. Estimating the maximum reserve capacity that a producer in a given price area can expect to sell at fixed price seems to be a challenging task. In the case of primary reserves (FCR), the maximal volume that can be supplied is limited by the droop setting for the generators. Currently, a minimum droop setting of 2 % is allowed, corresponding to a maximum reserve capacity of 10% of the installed capacity. In the case of secondary (FRR-A) and tertiary (FRR-M) reserves, the technical limitations are not that tight. If considering sales to these markets, it seems more realistic to define the maximum capacity limit based on observations. The Nordic TSOs request reserve capacity for the synchronous system so that pre-defined requirements for the different types of reserve products are satisfied. Knowing these system-wide constraints, one could in principle use a fundamental market model (such as the EMPS model) to forecast the both the prices and volumes requested by the TSO for each price zone. These price/volume scenarios will then be correlated with the energy price scenarios obtained from the same model.

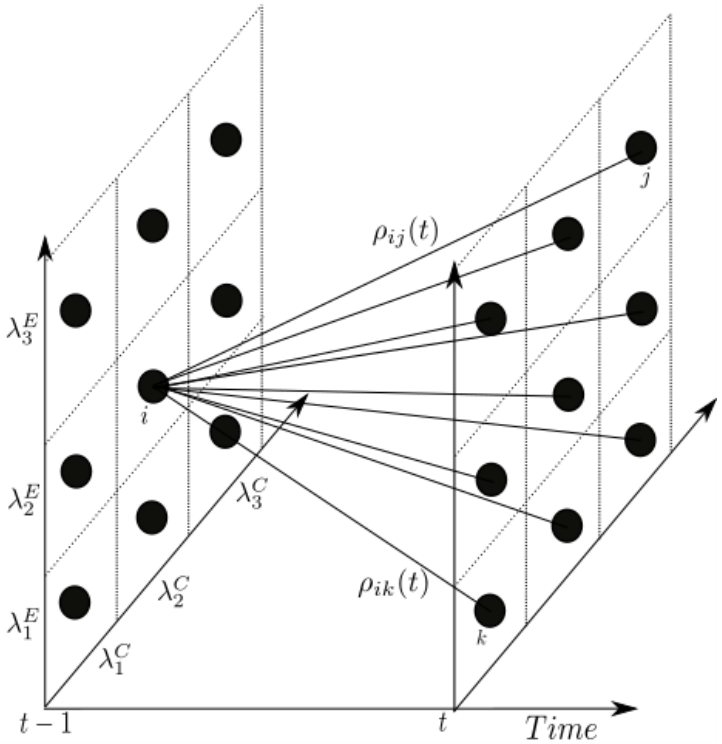

Figure 4 Treatment of two price processes in the combined SDP/SDDP scheme. λ^E is the energy price and λ^C the reserve capacity price in a given node.

2.5 The weekly decision problem

Once we know the decision stage (week no.) and the realizations of stochastic variables, the weekly decision problem can be formulated as an LP problem. The problem will have the same structure (variables and constraints) in both the forward and backward iteration. The mathematical formulation is presented in the following. The nomenclature is provided in Appendix A, and is sporadically repeated in the text below.

The objective function is stated in (1.1), maximising revenues from the two markets. The first two terms describe the weekly revenue from the reserve capacity and the energy markets, respectively. Reserve

capacity is sold for the week ahead ($t+1$) while energy is sold for the current week (t). For simplicity of notation, the week index is only used to specify variables and parameters for the week ahead.

$$Z_t = \max\left\{\sum_{b\in B}\tau_b\varsigma^C_{b,t+1}\lambda^C_{b,t+1}c_{b,t+1} + \sum_{k\in K}\sum_{h\in H}\varsigma^E_k\lambda^E_p e_{kh} - \sum_{h\in H}\varphi w_h + \alpha_{p,t+1}\right\} \tag{1.1}$$

An auxiliary variable w_h is introduced in the third term in Eq. (1.1) allowing the model to artificially supply water to the reservoir at a high cost φ. This variable is needed to ensure that the LP problem can be solved for all possible combinations of initial reservoirs and inflows. Finally the variable $\alpha_{p,t+1}$ represents the future expected profit, and is constrained by cuts as will be discussed later.

Each hydro reservoir h has a water balance for each time step k, as described in Eq. (1.2). There is one variable associated with each of the possible water ways (D = discharge through turbine, S = spillage and B = bypass). The initial reservoir level enters as a parameter for the first time step within the week. The accumulated weekly inflow I_h is distributed among time steps according to their relative duration. The dual value π_{kh} associated with Eq. (1.2) is the water value for reservoir h in time step k.

$$v_{kh} + \sum_{s\in S_h} q^D_{khs} + q^S_{kh} + q^B_{kh} - w_h - \sum_{j\in\Omega_h}\left(\sum_{s\in S_j} q^D_{kjs} + q^S_{kj} + q^B_{kj}\right) = v_{k-1,h} + \tilde{\tau}_k I_h \quad (\pi_{kh}) \quad [\mathrm{Mm^3}] \tag{1.2}$$

Eq. (1.3) is included to keep track of the energy generation (and sales) resulting from the discharge. Water is discharged through the power station using one variable per discharge segment. These segments will be used in decreasing order according to their energy equivalent η_{hs}, provided that η_{hs} decreases with s (concave PQ curve). As long as e_{kh} has a high enough upper bound so that energy is sold to a market with unlimited capacity, the dual value of (1.3) equals the energy price $\varsigma^E_k\lambda^E_p$ for the time step in question.

$$e_{kh} - \sum_{s\in S_h}\eta_{hs} q^D_{khs} = 0 \qquad [\mathrm{MWh}] \tag{1.3}$$

So far, the presented constraints are well known for any hydropower scheduling model representing the energy market only. Next we will look at the constraints concerning reserve capacity sales and allocation. The reserve capacity $c_{b,t}$ sold in the previous week enters the optimization problem as a parameter. It therefore becomes a commitment that should be met by allocating reserve capacity on the individual power stations in the water course according to Eq. (1.4). The dual value μ_b can be interpreted as the cost (or lost profit) when increasing the reserve capacity requirement with one MW.

$$\sum_{h\in H} r_{kh} = c_{b,t} \quad (\mu_b) \qquad [\mathrm{MW}] \tag{1.4}$$

We require that the reserve capacity allocated at a power station should be symmetric and spinning. That is, the allocated capacity should be available for both upwards and downwards regulation, and the station should be running at a power output that allows for down regulation. The latter is included in Eq. (1.5). In case the scaling factor γ_h equals 1, Eq. (1.5) states that if a reserve capacity r_{kh} is allocated a station h in time step k, the station should at least generate energy at that level. This modelling may lead to unfortunate generation schedules as illustrated by the following example. Consider a station with minimum power output at 50 MW and a maximum reserve capacity delivery of 33 MW. Being a linear model, we cannot force the operation of the station to either zero or above 50 MW. In cases where reserve capacity prices are high and

water values are higher than power prices, the model wants to sell reserves but not energy, and may therefore see 33 MW as the optimal operating point. The gamma factor is included in Eq. (1.5) to make operation below a defined minimum power level less profitable. We define $\gamma_h = \max\left\{\frac{P_h^{\min}}{R_h^{\max}}, 1.0\right\}$, in the example above $\gamma_h = \frac{50}{33}$. The scaling factor will therefore function as a means to punish artificially low energy generation for the sole purpose of delivering capacity reserves. In physical operation, down regulation should be available from the operation point to the minimum power output. We suggest that this is modelled by increasing the scaling factor, e.g. so that $\gamma_h = \frac{P_h^{\min} + R_h^{\max}}{R_h^{\max}}$.

$$\gamma_h r_{kh} - \frac{1}{\tau_k} e_{kh} \leq 0 \quad \left(\kappa_{kh}^{-}\right) \quad [\text{MW}] \tag{1.5}$$

Eq (1.6) ensures that the station has room for upward regulation.

$$r_{kh} + \frac{1}{\tau_k} e_{kh} \leq P_h^{\max} \quad \left(\kappa_{kh}^{+}\right) \quad [\text{MW}] \tag{1.6}$$

Eq. (1.7) ensures that there is enough water in the reservoir to deliver up-regulation reserves at the lowest efficiency η_{hS} for the entire time period in question. In the case of primary and secondary regulation reserves, this constraint may seem conservative, as the activation of these reserves will not span several consecutive hours. On the other hand, since we have not included time delay and ramping constraints when modelling the watercourse, Eq. (1.7) will translate to a water requirement for the reservoir in question plus all upstream reservoirs.

$$v_{kh} - \frac{\tau_k}{\eta_{hS}} r_{kh} \geq V_{kh}^{\min} \tag{1.7}$$

The expected future profit α is constrained by cuts as in Eq. (1.8). These cuts are built in the backward iteration and stored per time stage and price node. All state variables should be represented in the cuts. In our case the state variables are the reservoir levels at the end of the week, the weekly inflow in the current week and the reserve capacity sold per block for the next week. We have omitted the treatment of inflow from the cuts in Eq. (1.8) for simplicity, see e.g. [1] for more details.

$$\alpha_{p,t+1} - \sum_{h \in H} \pi_{phl} v_{kh} - \sum_{b \in B} \mu_{pbl} c_{b,t+1} \leq \beta_{pl} \tag{1.8}$$

The cuts are created based on the dual values π and μ from Eq. (1.2) and Eq. (1.4) obtained in the backward iteration. These values have different signs; π is positive since it gives the increase in profit when increasing the right-hand side in Eq (1.2) (initial reservoir or inflow) marginally, and μ is negative since it gives the marginal decrease in profit (or increase in operational cost) when increasing the right-hand side in Eq. (1.4) (reserve capacity sold in previous week). Figure 5 illustrates how the future expected profit function is constrained as a function of a given reservoir level (left) and sold reserve capacity in a given block (right).

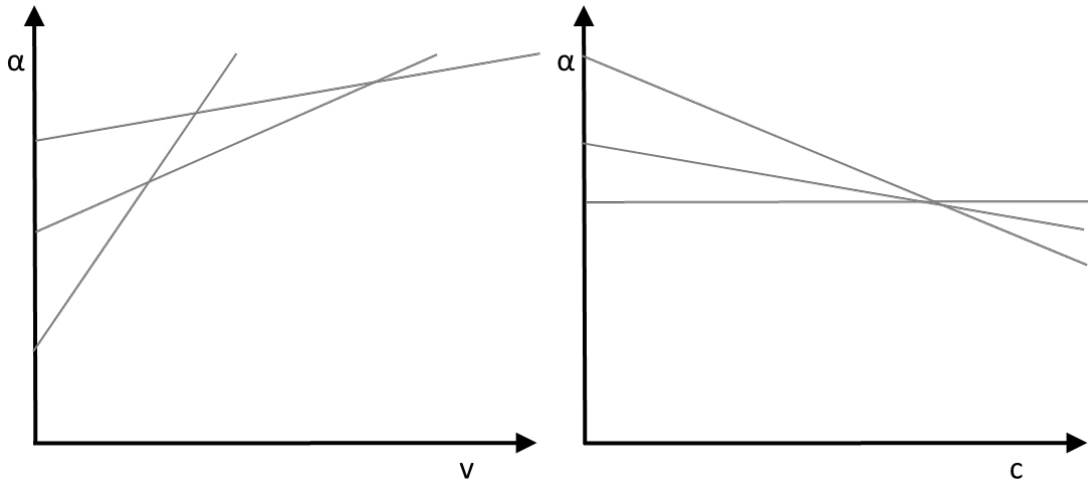

Figure 5 Illustration of expected future profit as a function of a reservoir volume (left) and a sold reserve capacity (right).

The number of cuts stored increases with the number of iterations, and normally only a few of these cuts are binding when solving a specific LP problem. For computational efficiency, cuts are therefore treated by relaxation. First we solve the LP problem without any cuts. Subsequently search through the list of cuts to find the most violated[5] and add that one to the LP problem before re-solving.

All variables are non-negative and have upper bounds. In particular it is worth pointing at the sold and allocated reserve capacity, shown in Eq. (1.9) and (1.10), respectively. As discussed in section 2.4, one can e.g. define the maximum reserve capacity volume R_h^{max} that a generator can deliver based on technical considerations or derive it from a regional requirement obtained from a fundamental market analyses.

$$0 \leq c_{b,t+1} \leq \sum_{h \in H} R_h^{\max} \tag{1.9}$$

$$0 \leq r_{kh} \leq R_h^{\max} \tag{1.10}$$

[5] A cut that is not included in optimization problem is violated if the inequality in (1.8) does not hold after inserting the solution from the optimization problem.

3 Impact on water values

We will in the following simplify the formulation of the weekly decision problem provided in section 2.5 in order to analytically evaluate the impact of capacity sales on the water values.

3.1 Analytical expressions

We define a system comprising a single hydropower reservoir with a power station selling both energy and reserve capacity at a pre-defined price for one time step. As a further simplification we assume that the reservoir will not hit its maximum or minimum levels in this time step, but there is an upper bound Q^{max} (which equals P^{max}) on discharge. The conversion factor between water and energy is assumed 1.0, so that water, energy and capacity refer to the same unit. The simplified problem formulation defined by equations (1.11)-(1.16) below. Capacity sold in this week to be delivered the next week is denoted c_{t+1}, and capacity sold the previous week entering this week as an obligation is denoted C_t.

$$Z_t = \max\left\{\lambda^C c_{t+1} + \lambda^E q_t + \alpha_{t+1}\right\} \tag{1.11}$$

$$v_t + q_t = v_{t-1} + I_t \quad (\pi) \tag{1.12}$$

$$q_t \geq C_t \quad (\kappa^-) \tag{1.13}$$

$$q_t \leq Q^{\max} - C_t \quad (\kappa^+) \tag{1.14}$$

$$\alpha_{k,t+1} - \tilde{\pi}_k v_t - \tilde{\mu}_k c_{t+1} \leq \tilde{\beta}_{k,t} \quad , k = 1...NK \quad (\psi_k) \tag{1.15}$$

$$0 \leq c_{t+1} \leq C^{\max} \quad (\chi^-, \chi^+) \tag{1.16}$$

Next we define the Lagrange function for the problem:

$$\begin{aligned} &L(c_{t+1}, q_t, \alpha_{t+1}, v_t, \pi, \kappa^-, \kappa^+, \psi_k, \chi^-, \chi^+) = \\ &\quad \lambda^C c_{t+1} + \lambda^E q_t + \alpha_{t+1} \\ &+ \pi\left(v_{t-1} + I_t - v_t - q_t\right) \\ &+ \kappa^-\left(C_t - q_t\right) + \kappa^+\left(Q^{\max} - C_t - q_t\right) \\ &+ \sum_{k=1}^{NK} \psi_k\left(\tilde{\beta}_{k,t} - \alpha_{t+1} + \tilde{\pi}_k v_t + \tilde{\mu}_k c_{t+1}\right) \\ &+ \chi^-\left(0 - c_{t+1}\right) + \chi^+\left(C^{\max} - c_{t+1}\right) \end{aligned} \tag{1.17}$$

After differentiating the Lagrange function with respect to each of the four variables we can obtain the following Kuhn-Tucker conditions:

$$\frac{\partial L}{\partial c_{t+1}} = \lambda^C + \sum_{k=1}^{NK} \psi_k \tilde{\mu}_k - \chi^- - \chi^+ = 0 \tag{1.18}$$

$$\frac{\partial L}{\partial q_t} = \lambda^E - \pi - \kappa^- - \kappa^+ = 0 \tag{1.19}$$

$$\frac{\partial L}{\partial \alpha_{t+1}} = 1 - \sum_{k=1}^{NK} \psi_k = 0 \tag{1.20}$$

$$\frac{\partial L}{\partial v_t} = -\pi + \sum_{k=1}^{NK} \psi_k \tilde{\pi}_k = 0 \tag{1.21}$$

If we assume that only one cut *k* is binding, we see from (1.20) that $\psi_k = 1$, and from (1.21) that $\pi = \tilde{\pi}_k$. Consequently, the water value equals the cut coefficient of the binding cut *k*.

By joining (1.19) and (1.21) we obtain the following expression:

$$\pi = \sum_{k=1}^{NK} \psi_k \tilde{\pi}_k = \lambda^E - \kappa^- - \kappa^+ \tag{1.22}$$

Expression (1.22) tells us that if there are no constraints on discharge, the water value will equal the energy price. If discharge is constrained by the lower bound in (1.13) for the purpose of keeping the unit spinning, then $\kappa^- < 0$ and the water value will be higher than the energy price. Conversely, if discharge is constrained by the upper bound in (1.14), then $\kappa^+ > 0$ and the water value will be lower than the energy price.

3.2 The impact of reserve capacity commitment – An example

Consider the following example. A reservoir has a capacity of 20, an initial level of 10, inflow in the current decision period of 3, a maximum discharge capacity of 4, and a symmetric reserve capacity commitment of 2. The energy price is 4 and there are two cuts approximating the future expected profit, with coefficients 5 and 3, as illustrated in Figure 6.

If we ignore the reserve capacity commitment, the optimal decision is to produce the inflow, and end up at a reservoir level of 10 and a water value of 4 (equal to the energy price). In this case, both cuts will be binding, as illustrated by the red circle in Figure 6.

If we treat the reserve capacity commitment, equation (1.14) tells us that the upper discharge boundary is 4-2 = 2. The optimal decision is now to produce 2, ending up at a reservoir level of 11. As illustrated by the blue dot in Figure 6, cut number 2 is binding and the corresponding water value is 3. According to equation (1.22) $\kappa^+ = 1$.

We find that the requirement for available up-regulation capacity is binding and contributes to a reduced water value.

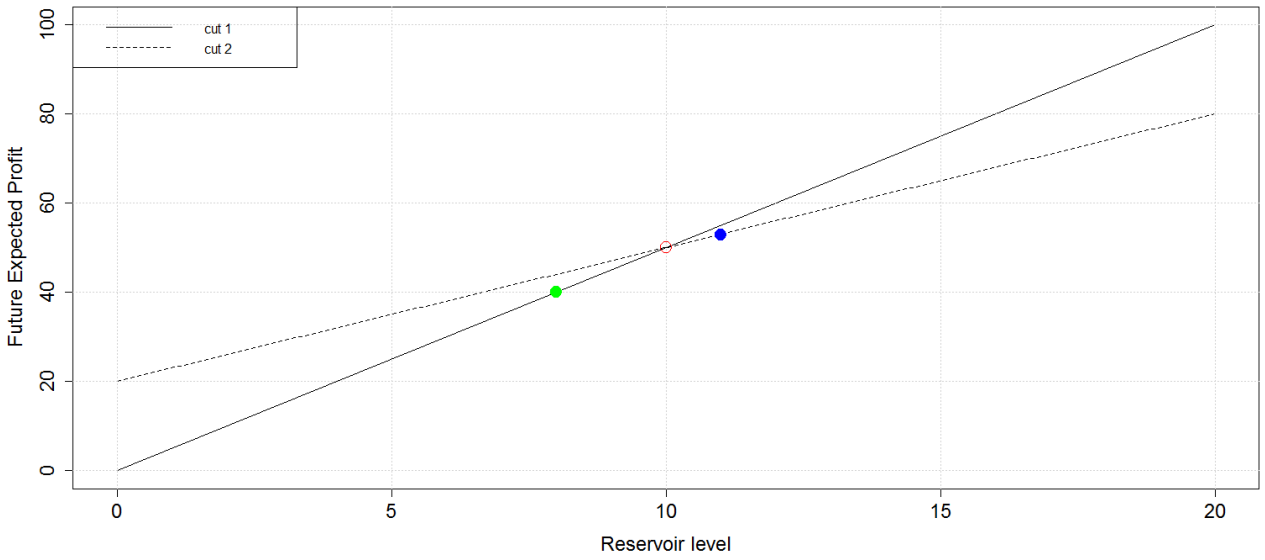

Figure 6 Cuts constraining future profit.

Starting at a different initial point, with a reservoir of 7 and the same characteristic data as before, the optimal decision would be to store the inflow and end up at a reservoir level of 10, with a water value of 4. Again this corresponds to the red circle in Figure 6. However, due to the reserve capacity commitment, we need to run the station and can only store 1 unit, ending up at a reservoir level of 8. This time cut number one is binding and the water value is 5. According to equation (1.22) $\kappa^- = -1$.

We find that the requirement for available down-regulation capacity is binding and contributes to an increased water value.

3.3 The impact of week-ahead reserve capacity sales

The decision whether or not to sell week-ahead reserve capacity is essentially an evaluation of the prevailing reserve capacity price against the expected week-ahead cost (or loss of profit) of delivering reserves.

By reformulating Eq. (1.18), we obtain:

$$-\sum_{k=1}^{NK} \psi_k \tilde{\mu}_k + \chi^- + \chi^+ = \lambda^C \tag{1.23}$$

At the optimal solution, there is a balance between the binding cut coefficient(s) and the reserve capacity price. If reserve capacity sales hits either of its boundaries in Eq. (1.16), the reduced costs will contribute as well. Note that the coefficients $\tilde{\mu}_k$ are negative indicating the reduced future expected profit associated with the sales of an additional unit reserve capacity.

From the optimality condition in (1.23) it can be seen that the prevailing price for reserve capacity can impact which cut(s) that are binding and therefore also impact the water value.

In case sales of reserve capacity and energy are carried out simultaneously for the same decision stage (co-optimized), the problem formulated in section 3.1 is slightly reformulated, making c_t a variable that enters the objective function and removing c_{t+1}. Thus, c_{t+1} will not be included in the cut, and we need to consider the impact of reserve sales in the current period rather than the next. The reformulation of the adjusted version of equation (1.23) will take the following form:

$$-\kappa^- + \kappa^+ + \chi^- + \chi^+ = \lambda^C \tag{1.24}$$

The difference between the two modelling approaches can be explained as follows. In the sequential approach we sell capacity for the next decision stage with uncertain information about energy price and inflows for the next decision period. This uncertainty is described by discretized probability functions, and is reflected in the cuts of type (1.15). When considering the capacity commitment from the previous week, cut coefficients are computed as the probability-weighted contributions of κ^- and κ^+ from (1.13) and (1.14) for all possible realizations of uncertain variables. For the co-optimization approach, we sell energy and capacity simultaneously with perfect knowledge of all stochastic variables within the week. Thus, the sales of reserve capacity can be fitted to match the energy sales for the given realization of stochastic variables, and the lost revenue from the energy market is then likely to be underestimated compared to the sequential approach.

3.4 Reserve cost curve

Reserve capacity can to a certain extent in some water courses be sold as a by-product at little additional cost. This is e.g. the case for a generator running at best point with available capacity for both up- and down regulation. There will however be threshold after which further sales of capacity more severely impacts system operation, and therefore is associated with a higher cost. This is the case whenever sales of capacity leads to a different optimal solution than would otherwise have been found in the energy-only case. In such cases there is a lost revenue in the energy market (representing the opportunity cost).

Consider the cuts describing the expected future profit function restated in Eq. (1.25). If we assume that capacity is reserved for one block and that reservoir levels takes a predefined level, expected future profit for this decision stage can be presented as a function of the committed capacity for the next decision stage, as shown to the left in Figure 7. This function is concave and non-increasing.

$$\alpha_{p,t+1} - \sum_{h \in H} \pi_{phl} v_{kh} - \sum_{b \in B} \mu_{pbl} c_{b,t+1} \leq \beta_{pl} \tag{1.25}$$

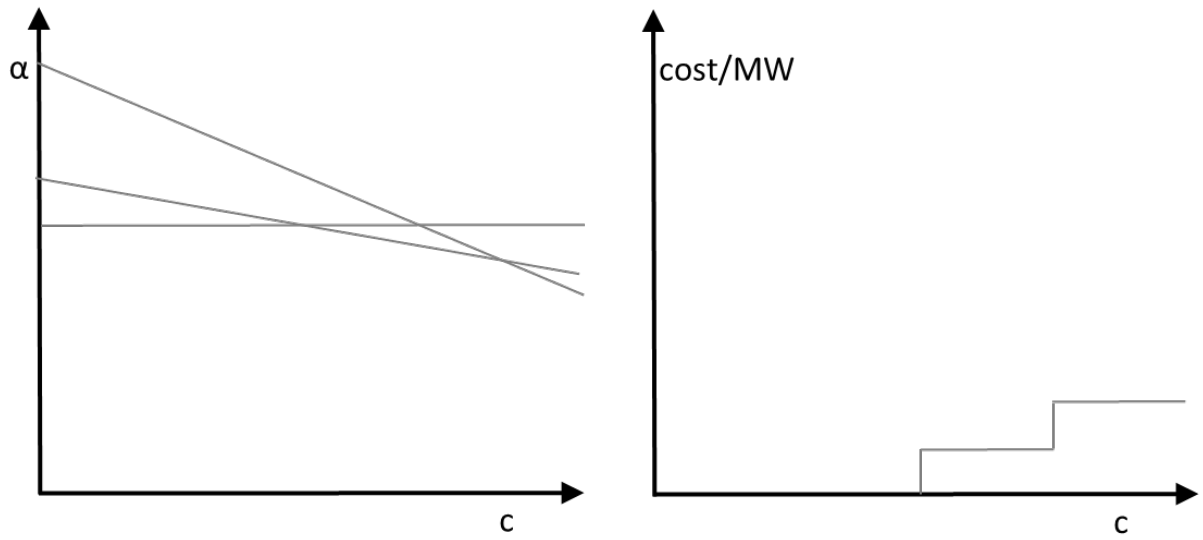

Figure 7 Illustration of expected future profit (left) and marginal cost (right) as a function of sold reserve capacity for the next decision stage.

The function in Figure 7 provides information on the cost of delivering reserves, given a pre-defined price node and a combination of reservoir levels. It can provide useful information and decision aid for pricing of reserve capacity. The coefficient of the binding cut for a given reserve capacity level (left in the figure) is the marginal cost of delivering that capacity (right in the figure).

The process of selling capacity is in principle a two-stage optimization problem embedded in the overall multi-stage hydro scheduling problem. The income from the capacity reserve sold in the current week is evaluated against the expected cost of meeting that obligation the next week. Thus, the marginal cost curve shown to the right in Figure 7 does not explicitly depend on the reserve capacity price in that decision stage. It will however implicitly depend on the reserve capacity price-level, since this level will impact the use of water and therefore also the water values.

4 Case studies

In this section we present analyses performed with the presented model on a simple test system comprising one hydro "module" (reservoir and station). The system is kept simple in order to facilitate transparent result interpretation and discussions to the following questions:

- How are water values affected when selling both energy and reserve capacity?
- How do different modelling assumptions and system configurations impact the previous answer?
- Can we use the model to calculate the cost of delivering reserves?

4.1 System description

The system comprises one hydro module (reservoir and power station). The reservoir has a storage capacity of 150 Mm3 and an average annual regulated inflow of 450 Mm3, giving a degree of regulation of 0.33. The stations PQ-curve is shown in Table 1. When running at its best point, the station can sell 30 MW of reserve capacity.

Table 1 PQ description for hydropower station.

Point	P [MW]	Q [m3/s]	Efficiency [MW/m3/s]
1	80	20.00	4.00
2	90	22.56	3.91
3	100	25.16	3.85
4	110	27.79	3.79

The energy price series was obtained from the EMPS model. Based on these series we generated a price model with two weekly average energy price nodes, as shown in Figure 8. A transition probability matrix was computed describing the probability of going from a given node in a given week to any of the nodes in the next week. The energy price follows a pre-defined (deterministic) profile within the week, scaled according to the average weekly price.

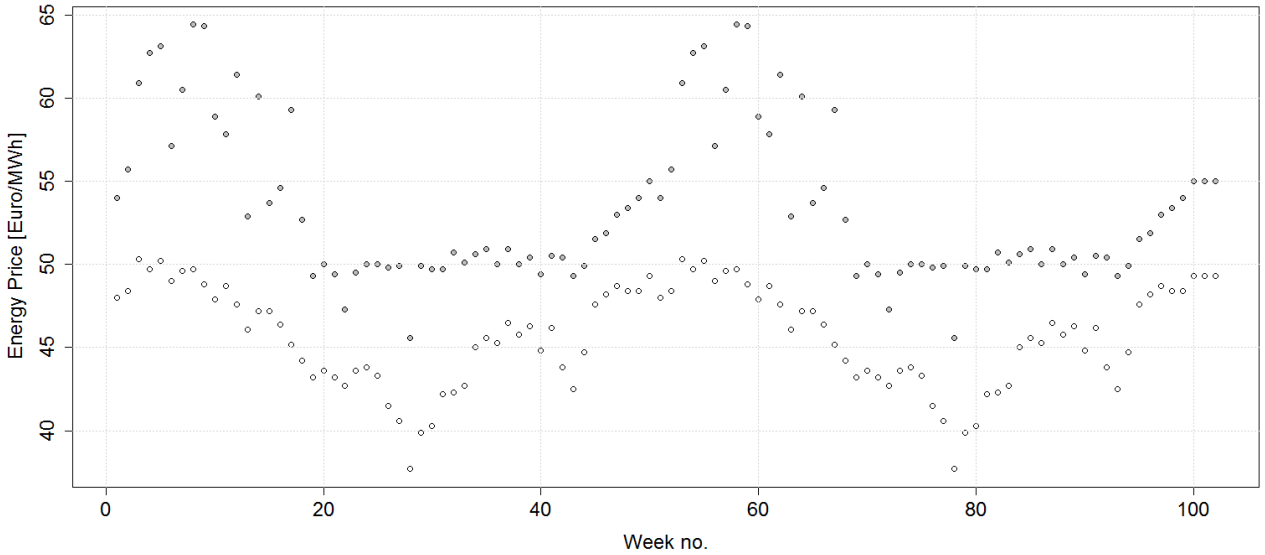

Figure 8 Energy price nodes representing two probable outcomes of average weekly energy price.

System operation was computed for a period of 104 weeks, considering 21 sequentially treated time steps within the week. Energy is sold per time step. Moreover, we let the model sell capacity for the week ahead in blocks, considering 3 blocks per week covering night (00:00-08:00), day (08:00-20:00) and evening (20:00-24:00). Note that the model does not consider activation of reserve capacity.

In sections 4.2 and 4.3 we present two different cases. In the first case, presented in section 4.2, we exaggerate the prices and volumes for reserve capacity in order to clearly point out the differences in scheduling and water values when considering sales of reserve capacity. In the second case, presented in section 4.3, we use observed reserve capacity prices obtained from the Norwegian FCR market. This case provides more realistic estimates on how the scheduling and water values are expected to change when considering sales of reserve capacity.

Each of the two cases were simulated in two different modes:

- **One-market mode (base case).** Only allows sales of energy. This case was prepared by setting the reserve capacity price to a negative value to prohibit sales of capacity.
- **Two-market mode.** Allows selling both energy and capacity following the model description previously defined.

The model was run for 20 main iterations considering 50 samples of inflow and energy price nodes in the forward iterations and 4 "openings" for inflow in the backward iteration. We used a set of end-value cuts to account for the end-of-horizon value of water and reserve capacity.

A convergence plot is shown in Figure 9, showing that the cost gap closes according to theory. We will not emphasize on the computational performance of the model in the following discussions.

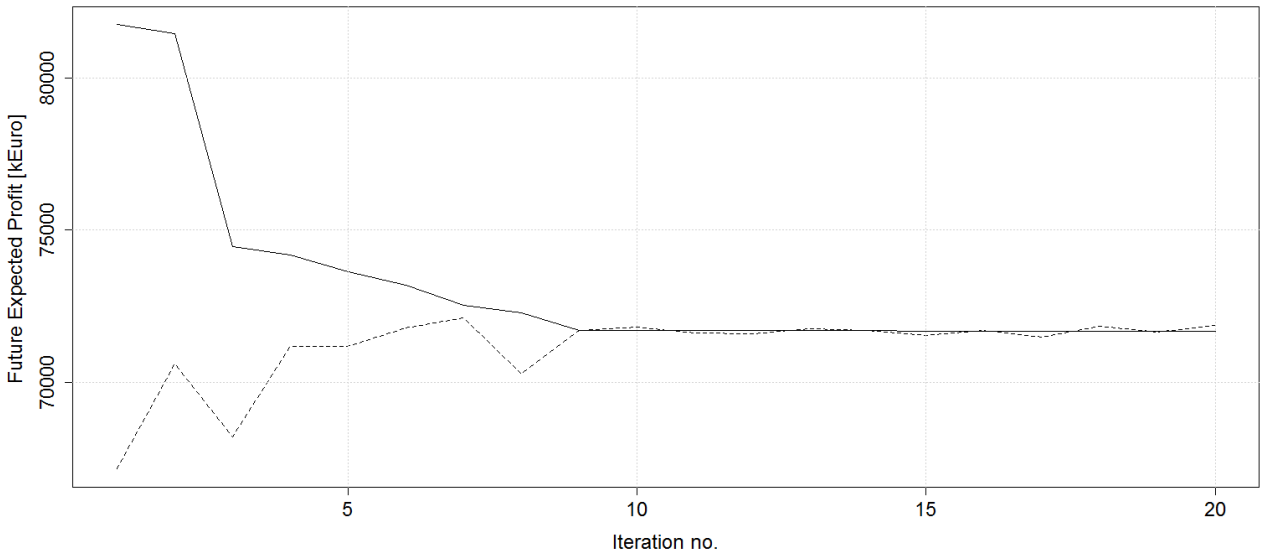

Figure 9 SDDP cost gap. Upper (solid) and lower (stapled) bound per iteration.

4.2 Case 1 – Exaggerated sales of reserve capacity

In this case we exaggerate the potential for selling capacity in order to provoke significant differences in the two scheduling modes. In the two-market mode we allow selling 44 MW of reserve capacity at a fixed price of 100 Euro/MW. Due to the extremely high reserve capacity price, the model will sell capacity at the maximum limit continuously. At low energy prices, the generation will be at 44 MW, enough to keep the committed reserves spinning. At high energy prices, the station will generate 66 MW leaving 44 MW for upward regulation. There is no minimum power production requirement and start-up cost for the station.

4.2.1 Reservoir operation

The reservoir operations for the two modes are shown in Figure 10 and Figure 11. Operation in the two-market mode is clearly less flexible than in the one-market mode, due to the operational constraints imposed by sales of capacity, as discussed above.

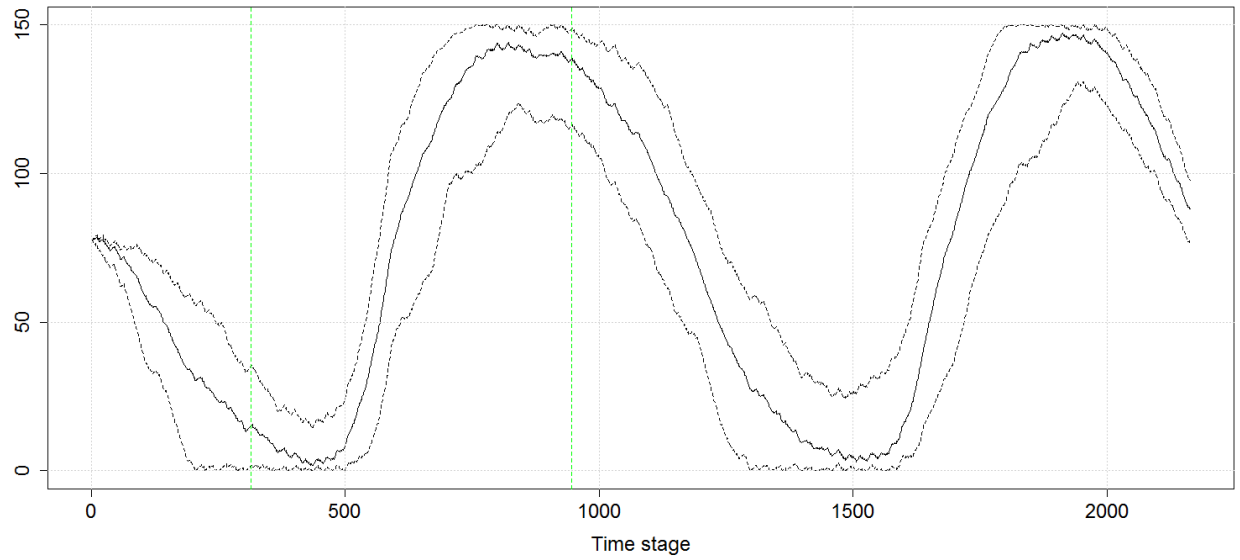

Figure 10 Simulated reservoir trajectories (max, mean and min) for the one-market mode, in Mm3. The green vertical lines mark weeks 15 and 45 for which water values are presented in section 4.2.2.

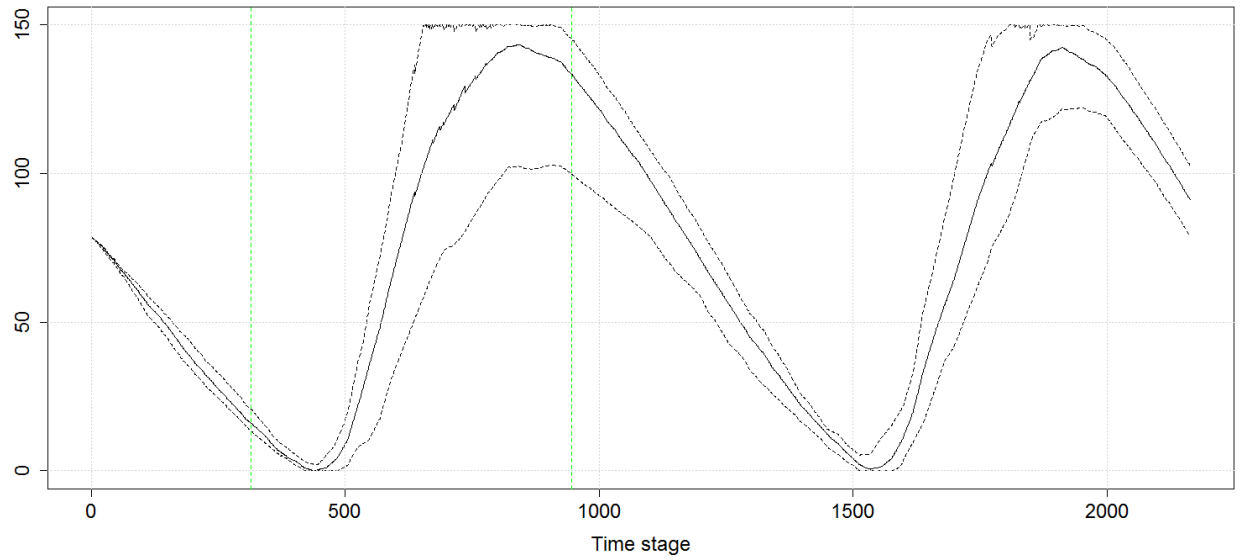

Figure 11 Simulated reservoir trajectories (max, mean and min) for the two-market mode, in Mm3. The green vertical lines mark weeks 15 and 45 for which water values are presented in section 4.2.2.

4.2.2 Water values

We will now study the water values computed in the two modes at different points in time. The water values are found by evaluating the cuts that have been generated in the backward iteration of the SDDP algorithm.

Although the presented case study only considers one reservoir, we describe a multi-reservoir approach below, as used in [6].

We first choose a reservoir that we would like to study results for, let's refer to this this as *hs*. We then evaluate, by assigning numerical values for the reservoir levels of all other reservoirs and capacity sales and moving these terms to the right-hand side. Each cut then becomes a linear inequality in one variable (v_{khs}).

$$\alpha_{p,t+1} - \pi_{phsl} v_{khs} \leq \beta_{pl} + \sum_{h \in H, h \neq hs} \pi_{phl} v_{kh} + \sum_{b \in B} \mu_{pbl} c_{b,t+1} \tag{1.26}$$

Next, we divide the reservoir variable v_{khs} into discrete levels, and move the contribution to the right-hand side in (1.27). The cut with the lowest numerical right-hand side value will be binding for that particular state. By repeating the same procedure for all discrete values of v_{khs} we obtain a set of different cuts that are binding. The coefficient π_{khs} of the binding cut is then treated as the water value for that particular reservoir level.

Figure 12 and Figure 13 show the water values obtained from week 15 and 45, respectively. These weeks are indicated by the green vertical lines in Figure 10 and Figure 11. Since the cuts have been generated for simulated reservoir states, the water value curve has finer resolution in between the band of simulated reservoir states. This is clearly shown in Figure 12, where the two-market water value is constant (100 kEUR/Mm3) for the lowest reservoir levels (between 0 and 7 % filling). 100 kEUR/Mm3 was set as the cost of buying artificial water to the system, and the binding cut with this coefficient has most likely been generated in an early iteration. The red lines in Figure 12 shows that the final simulated reservoir levels in week 15 are all higher than 11 %, so the constant shape of the water value curve for the two-market mode at low reservoir levels is most likely due to the fact that few cuts are generated for these levels.

Note that we have only stored and visualized simulated states for the final SDDP iteration. From the reservoir curves we can see that the two-market case has much smaller band in simulated trajectories than the one-market case in week 15. The red vertical lines in Figure 12 indicate the simulated extremal reservoir states in the two-market mode. For week 45 the band is not that different for the two modes, and is therefore not indicated in Figure 13.

In week 15 the reservoir is close to empty, and the water values obtained in the two-market mode are much more sensitive to changes in reservoir level than in the one-market mode. This is due to the limited flexibility when committing a majority of the capacity as reserve. If the reservoir level is high in week 15, we know that the station will not generate above 66 MW in the filling season due to high reserve capacity prices, and the risk of spillage is therefore high. Conversely, if the reservoir is low in week 15, we know that the station will at least run at 44 MW, and there is less water left for generating for the electricity market if prices are favourable.

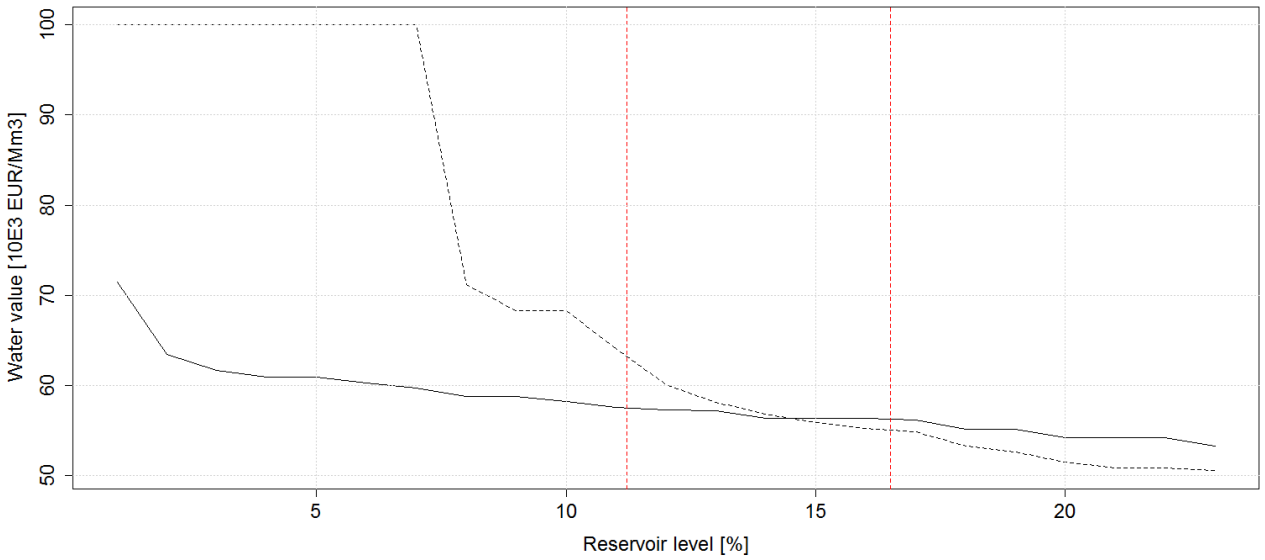

Figure 12 Water values obtained in the two modes for week 15. One-market values are solid-drawn and two-market values are stapled. The red vertical lines indicate the interval of simulated reservoir states in the final simulation of two-market mode.

In week 45 one enters the season where discharge normally is higher than inflow. Again we see from Figure 13 that the water values obtained in the two-market mode are significantly more sensitive to changes in reservoir than in the one-market mode. If the two-market mode reservoir level in week 45 is relatively low, the model is left with little flexibility in selling additional energy at favourable prices, and thus a high water value. On the other hand, if the reservoir level gets high enough, the maximum generation constraint of 66 MW (keeping 44 in reserve for upward regulation) will be frequently binding, and there is a higher risk of spilling next spring flood.

In summary we find that the water values obtained from the model can be explained by the theoretical results in section 3:

- A relatively low reservoir level and high reserve capacity prices implies that the down-regulation capacity constraint is binding, contributing to an increase in water value (compared to the one-market case)
- A relatively high reservoir level and high reserve capacity prices implies that the up-regulation capacity constraint is binding, contributing to a decrease in water value (compared to the one-market case)

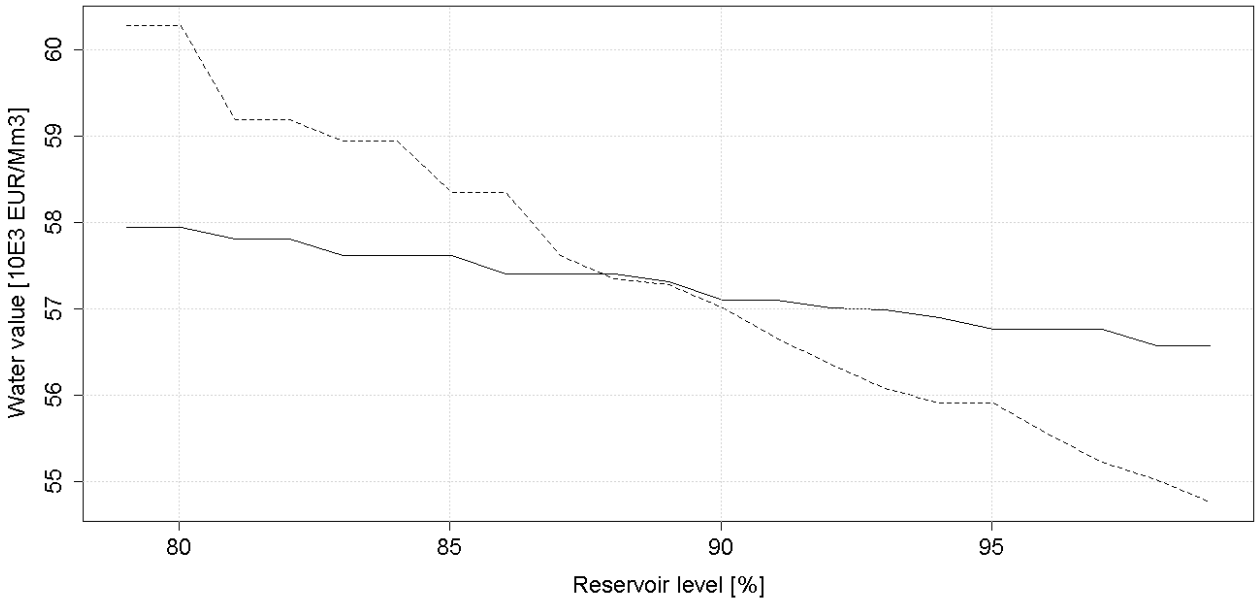

Figure 13 Water values obtained in the two modes for week 45. One-market values are solid-drawn and two-market values are stapled.

4.3 Case 2 – Modest sales of reserve capacity

In this case we take a more modest, but perhaps still optimistic, view on the potential for earning money in the reserve capacity market. We allow the system to sell at most 22 MW of spinning reserve capacity. If one relates this amount to the FCR market, it would correspond to a 1% droop setting. We assume a deterministic prices series comprising the FCR-N prices for 2013-2014 in NO2. The weekly average values are shown in Figure 14. There is no minimum power production requirement and start-up cost for the station.

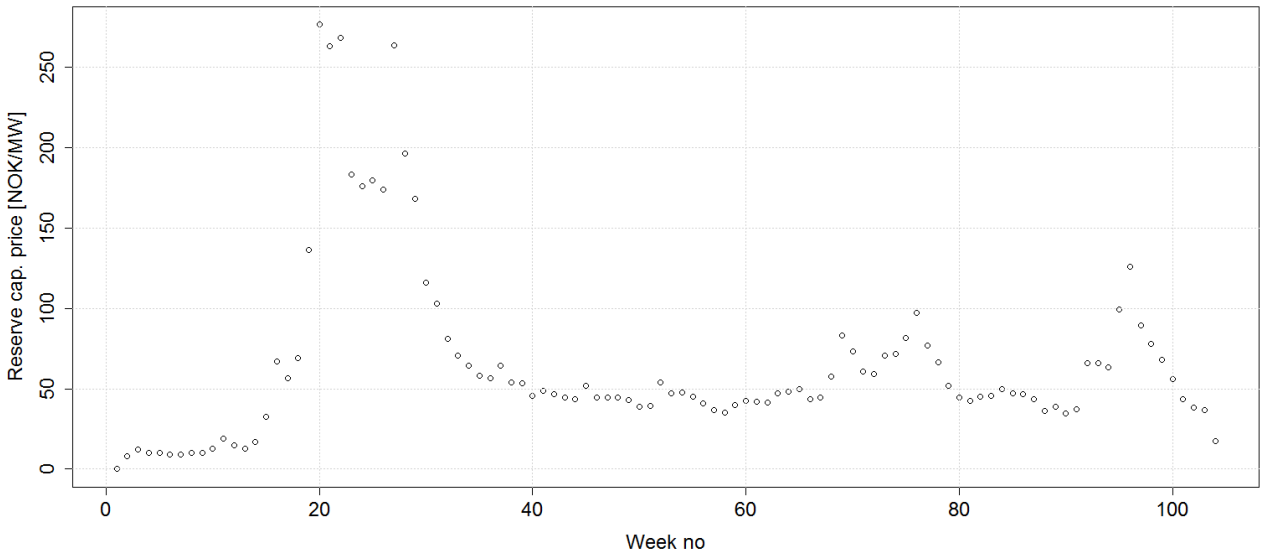

Figure 14 Reserve capacity prices used in case 2. Weekly average values for the FCR-N market in price area NO2.

4.3.1 Reservoir operation

The reservoir operation for the one-market mode was shown in Figure 10, and is shown in Figure 15 for the two-market mode. These figures indicate that the reservoir operation does not differ significantly.

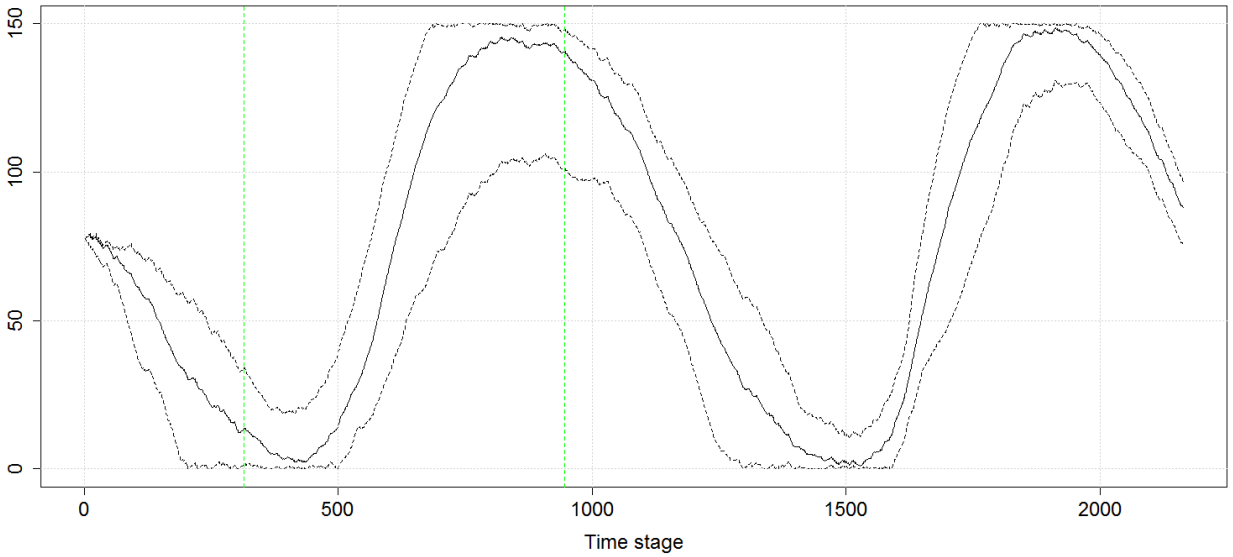

Figure 15 Simulated reservoir trajectories (max, mean and min), in Mm3, for the two-market mode in case 2.

The differences in schedules can be seen clearer when looking at the duration curves for generation, as shown in Figure 16. A significant part of the time the station is operated at either 88 MW or 22 MW in order to sell reserve capacity. The total amounts of energy produced in the two cases do not differ much.

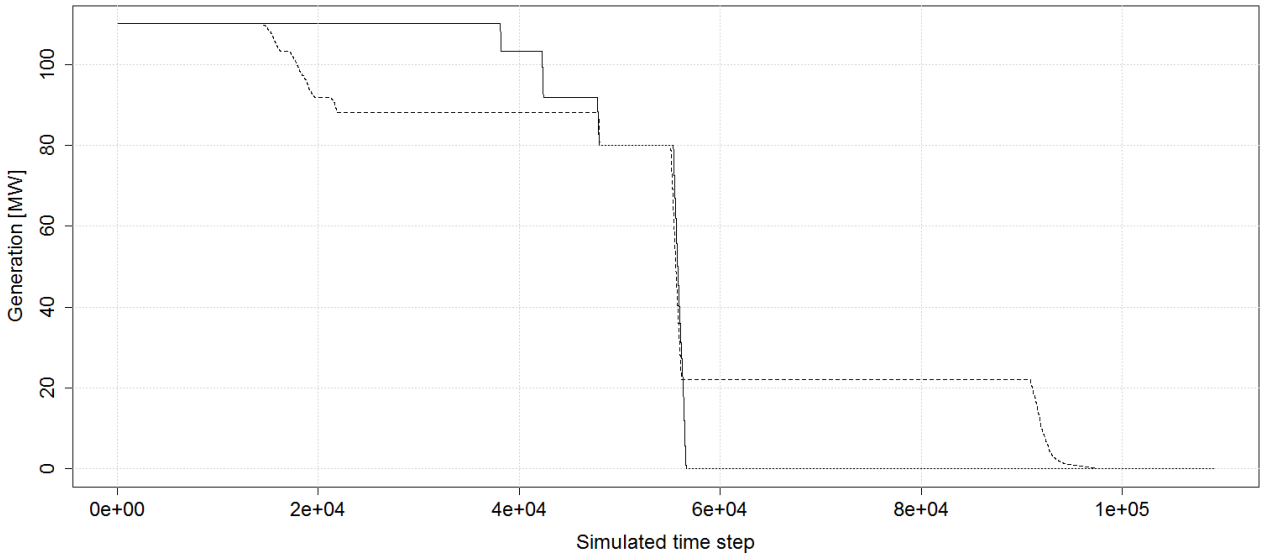

Figure 16 Duration curves for hydropower generation. Solid-drawn line for the one-market mode and stapled for the two-market mode.

The duration curve for generation for the two-market mode is re-drawn in Figure 17. The corresponding reserve capacity commitment (capacity sold in the previous week) is shown with a solid-drawn line. The chosen generation level can in most time steps be explained by the reserve capacity commitment. That is e.g. the case for the many time steps where generation is at 88 MW or 22 MW, for which there is a maximum reserve capacity commitment. Furthermore, the figure shows that operation around the best point (at 80 MW) is not directly tied to the capacity commitment.

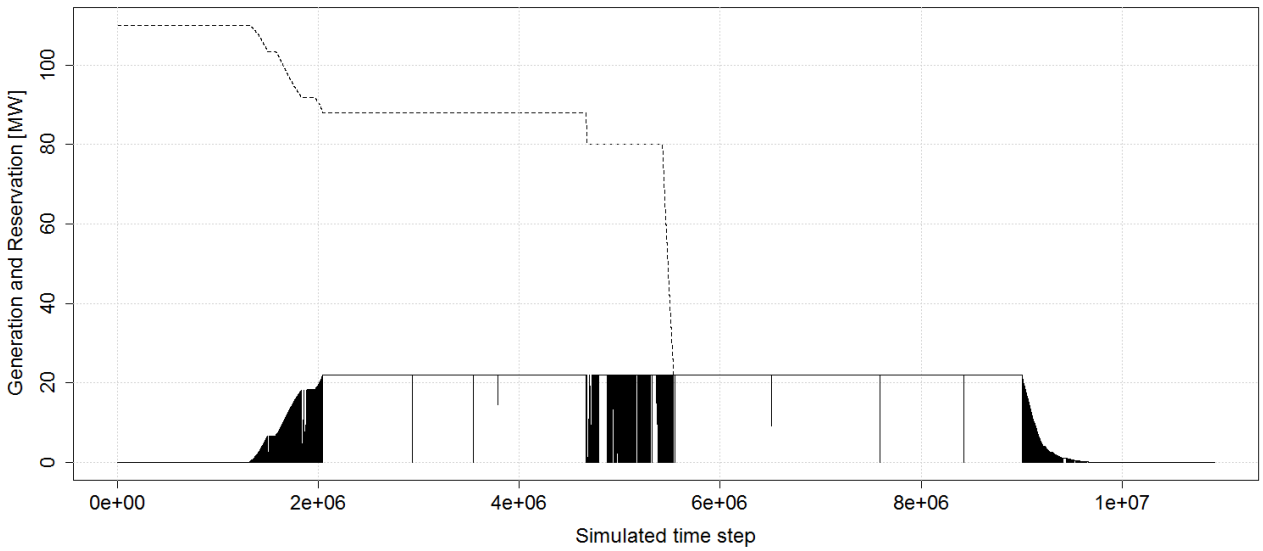

Figure 17 Duration curves for hydropower generation (stapled line) and corresponding sales of reserve capacity (solid-drawn line) for the two-market mode.

4.3.2 Water values

Figure 18 and Figure 19 show the water values obtained from week 15 and 45, respectively. As expected, the differences between the two modes are now less pronounced compared to case 1. In Figure 18, one can observe the same pattern as was shown for case 1, week 15 in Figure 12. The water values in the two-market mode are slightly higher at low reservoir levels and slightly lower at high levels. What seem to be less systematic crossing points of the two curves could be smoothed if we let the model run additional iterations and generate more cuts. This was confirmed by running some additional tests. For week 45, the water values are slightly higher in the two-market mode for all reservoir levels. This indicates that the operation from week 45 and forward is much more flexible than it was in case 1, shown in Figure 11. The risk of spilling is less dominating at high reservoir levels, so that the additional income from the reserve capacity market is reflected in the water values more or less independent of reservoir level in this week.

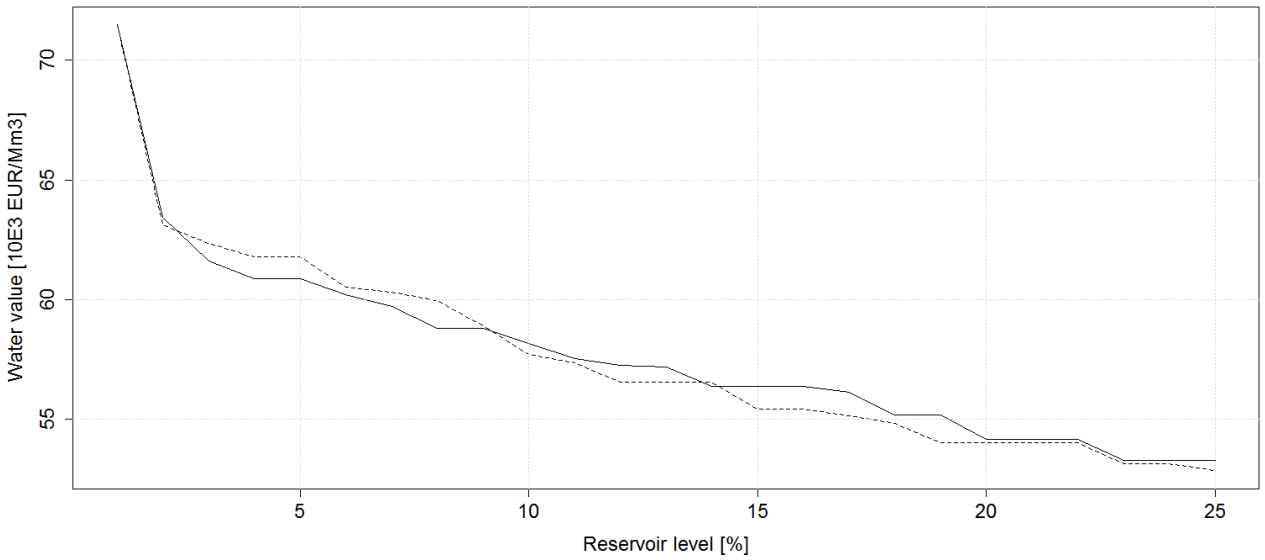

Figure 18 Water values obtained in the two modes for week 15 in case 2. One-market values are solid-drawn and two-market values are stapled.

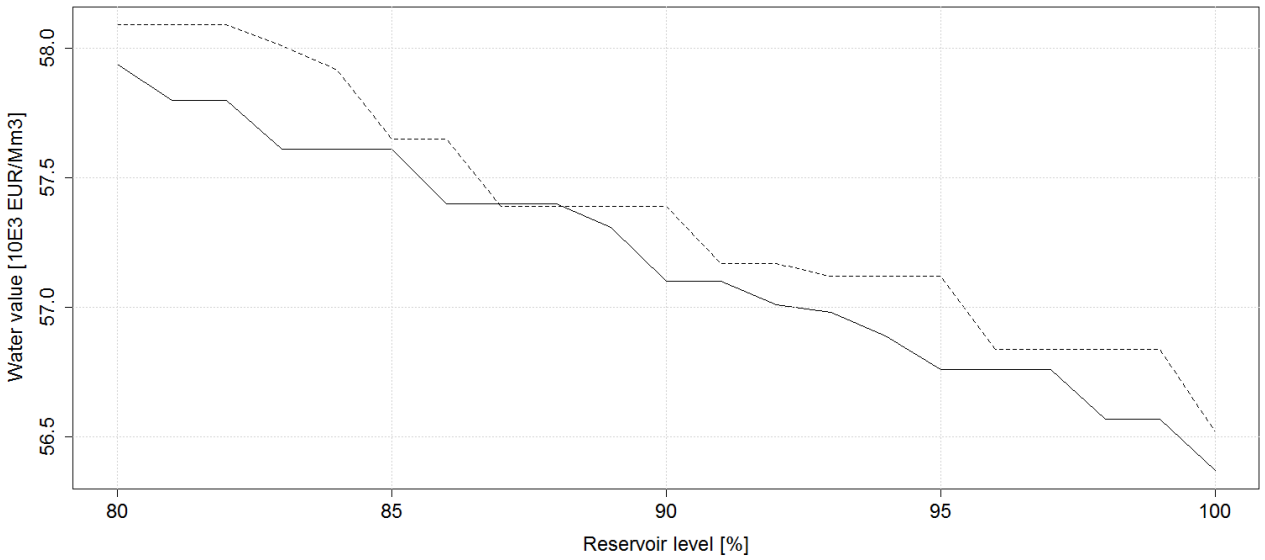

Figure 19 Water values obtained in the two modes for week 45 in case 2. One-market values are solid-drawn and two-market values are stapled.

4.3.3 Sensitivity analysis

We wanted to test how this water value comparison is affected by the utilization time of the power plants. To do so, we constructed two additional cases, one where the annual expected inflow was increased with 50%

(case *high*) and one where it was decreased by 50 % (case *low*). All other input data was as defined for case 2. Both the one- and two-market cases where tested in case *high* and case *low*.

Increasing the annual expected inflow in case *high* will lead to a higher utilization time for the station. Figure 20 and Figure 21 show the water values obtained in case *high* for week 15 and 45, respectively. Comparing the water values in week 15 in the base case (Figure 18) and case *high*, one can as expected see that they have decreased when increasing inflow. Comparing the one- and two-market values in each figure one sees that the two-market values are relatively lower than the one-market values for case *high*. This is also true for week 45 in Figure 21, which is opposite from what was observed in Figure 19. When increasing the inflow and thus the utilization time, the delivery of up-regulation reserves in the two-market case clearly contributes to lower the water values. Note that the expected income is always highest in the two-market mode, cf. section 4.3.4.

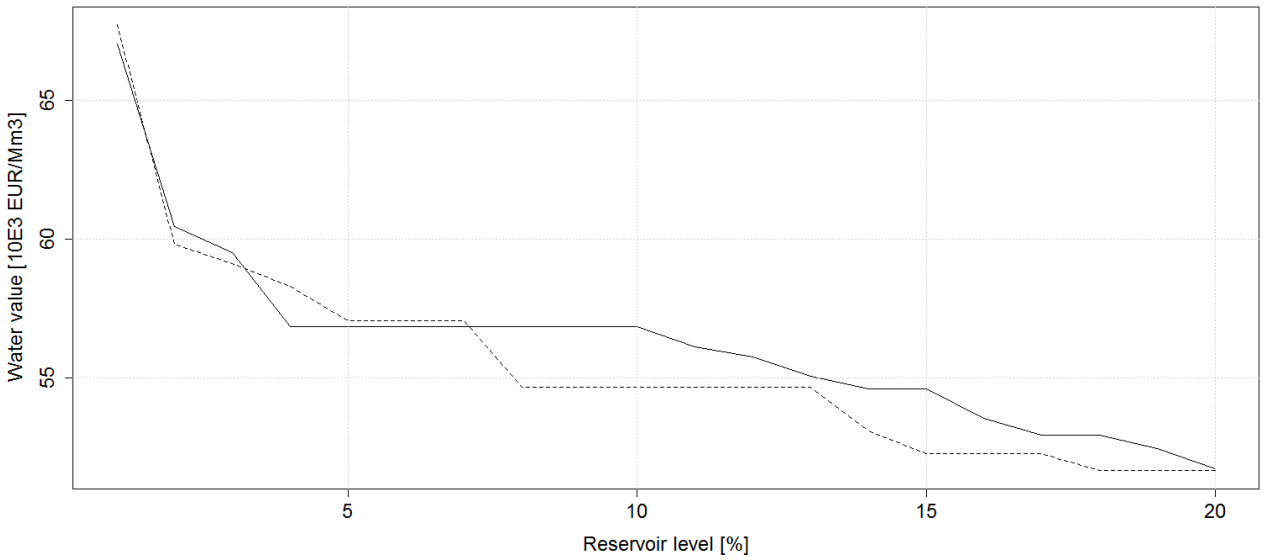

Figure 20 Water values obtained in the two modes for week 15 in case *high* (50% increase in inflow). One-market values are solid-drawn and two-market values are stapled.

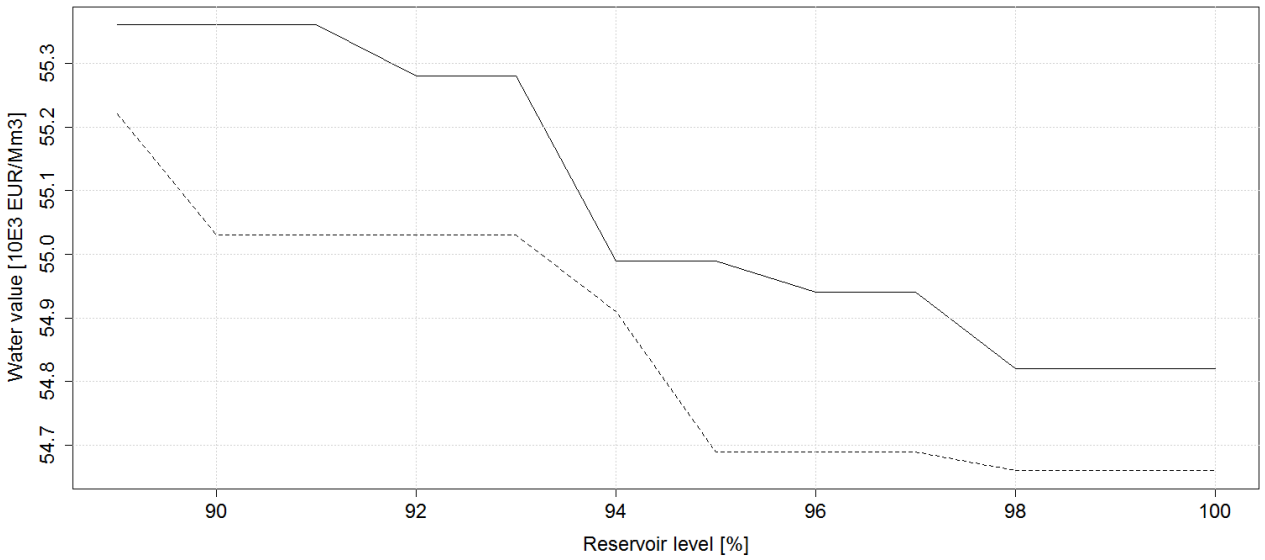

Figure 21 Water values obtained in the two modes for week 45 in case *high* (50% increase in inflow). One-market values are solid-drawn and two-market values are stapled.

Decreasing the annual expected inflow in case *low* will lead to a lower utilization time for the station. Figure 22 and Figure 23 show the water values obtained in case *low* for week 15 and 45, respectively. In both weeks, the water value is higher in the two-market case for the range of simulated reservoir levels along the x-axis. With a lower utilization time, the delivery of down-regulation reserves contributes to increase the water values.

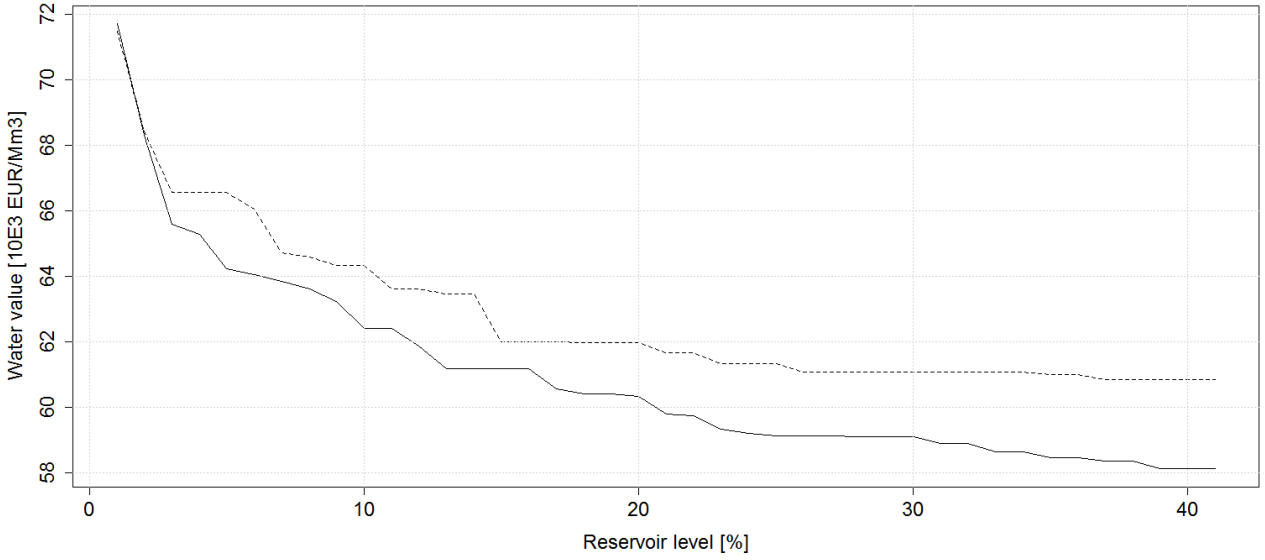

Figure 22 Water values obtained in the two modes for week 15 in case *low* (50% decrease in inflow). One-market values are solid-drawn and two-market values are stapled.

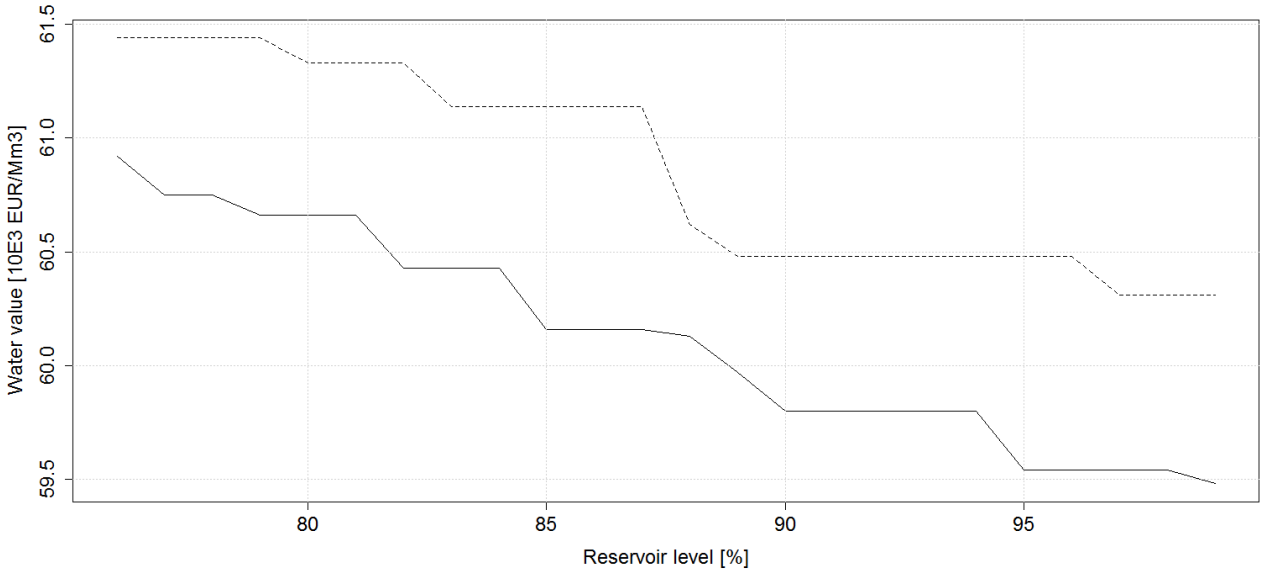

Figure 23 Water values obtained in the two modes for week 45 in case *low* (50% decrease in inflow). One-market values are solid-drawn and two-market values are stapled.

4.3.4 Profitability and cost of operation

After the strategy (cuts) has been computed, we performed a final simulation with 5000 samples of inflows and energy prices. This was done to obtain higher accuracy in the estimates of expected future profit. The large number of samples needed to obtain stable results is primarily due to the sampling in the price model. Recall that we are sampling one out of two price nodes (indicating either a high or low price) per week over the whole period of analyses. The coarse representation of energy price uncertainty calls for a larger number of samples than would be the case if the price model was more densely discretized with lower probabilities of jumping from a high to a low price.

The expected revenues from the two markets are shown in Table 2. When considering the revenues from both markets, sales of reserve capacity amounts 3.55 %. Since the two cases ends approximately at the same reservoir levels, we have not valuated the end reservoir in this comparison. Although sales of reserve capacity are expected to generate additional revenue of 1.97 MEUR, the profitability only increases by 1.63 MEUR. The difference of 0.34 MEUR represents the lost opportunities in the energy market when selling capacity.

Table 2 Expected revenues from the two markets.

Case	Energy sales [MEUR]	Capacity sales [MEUR]	Sum [MEUR]
One market	53.88	0.00	53.88
Two market	53.54	1.97	55.51

Altering the generation schedules from the one- to the two-market mode for the purpose of selling an extra product (reserve capacity) may come at a cost (lost opportunity). That cost can be evaluated by studying the

dual values κ^- and κ^+ from equations (1.5) and (1.6), respectively. κ^- will be less or equal to zero indicating the marginal decrease in profit when increasing the spinning requirement with one MW. κ^+ will be greater or equal to zero indicating the marginal increase in profit if increasing P^{max} with one MW. Note that increasing P^{max} will not allow for a higher energy generation since the discharge variables are already bounded; it only allows for more sales of reserves. If the dual values κ^- and κ^+ always are zero, selling capacity will not change the water values.

Up-regulation can be provided by the model at no cost if the station generates less than 88 MW in all time steps within a reserve block. If generating at least 88 MW in at least one of the time steps, the cost of up-regulation can be seen as $-\kappa^+$ (treating costs as negative). Conversely, down-regulation is provided at no cost if the station generates more than 22 MW in all time steps within a reserve block. If generating at most 22 MW in at least one of the time steps, the cost of down-regulation is κ^-. Table 3 shows the percentage of time in which there is a cost associated with selling reserves and the maximum and mean cost observed from the simulation results. Up-regulation reserves are primarily costly to deliver at winter time when energy prices are high and we want to run the power station at full output. Down-regulation reserves are primarily costly in the low-load season when the reservoir level is high and we otherwise would shut-down the station and save the water for higher energy prices.

Table 3 Cost of delivering capacity.

	Share of time [%]	Max [10E3/MW/ step]	Mean [10E3/MW/step]
Down-regulation	33.98	0.49	0.014
Up-regulation	37.40	0.84	0.025

4.3.5 Reserve cost curve

We can obtain the marginal cost of selling reserve capacity at different points in time by evaluating the cuts that have been generated in the backward iteration of the SDDP algorithm. This was discussed on a principle level in section 3.4.

We first choose a reserve block *bs* that we would like to study results for. We evaluate, by assigning numerical values for the reservoir levels and capacity sales in other time blocks than the one we are looking at (*bs*) and moving these terms to the right-hand side:

$$\alpha_{p,t+1} - \mu_{pbsl} c_{bs,t+1} \leq \beta_{pl} + \sum_{h \in H} \pi_{phl} v_{kh} + \sum_{b \in B, b \neq bs} \mu_{pbl} c_{b,t+1} \tag{1.27}$$

The expected future profit function then becomes a function of the reserve capacity sold in block *bs*. Next, we divide the capacity sales variable $c_{bs,t+1}$ into discrete levels. For each of these discrete levels we move the numerical value $\mu_{pbsl} c_{bs,t+1}$ to the right-hand side in (1.27). The cut with the lowest numerical right-hand side value will then be binding for that particular state. By repeating the same procedure for all discrete values of $c_{bs,t+1}$ we obtain a set of cuts that are binding for the different discrete reserve capacity levels. The coefficient μ_{pbl} of the binding cut is then treated as the marginal cost of selling a given level of reserve capacity.

A marginal cost curve for sales of capacity for reserve block 3 (20:00-24:00) in week 15 is shown in Figure 24. Since the future profit functions are concave and decreasing, we know that the shape of this curve should be non-decreasing. That is, the more capacity we sell, the more expensive it is for the system to deliver this

capacity. The rather high marginal cost at low capacity levels in Figure 24 may seem somewhat unexpected. One could expect that the first MWs sold do not involve any costs, provided that the station is running at its best point. The high values can be explained by the following moments:

- Capacity is sold per block, and one block covers a set of hours within each day. In this example, capacity is sold in three different blocks, while energy is sold in 21 sequential time steps. If we relax the requirement coupling time-steps to blocks when selling capacity, the marginal cost curve would in general be steeper than in Figure 24, starting at lower costs for little capacity.
- Capacity is sold for the week-ahead based on an expectation of energy price and inflow for the next week. Particularly the energy-price uncertainty will impact the shape of the marginal cost curve. If the energy price was known when selling reserves, the marginal cost curve would more often start at a lower level.

Moreover, the final (and tight/binding) cuts have been computed for a narrow band of sold capacity close to the maximum value. Thus, one should not expect that the cost profile in Figure 24 accurately cover the complete range of capacity values.

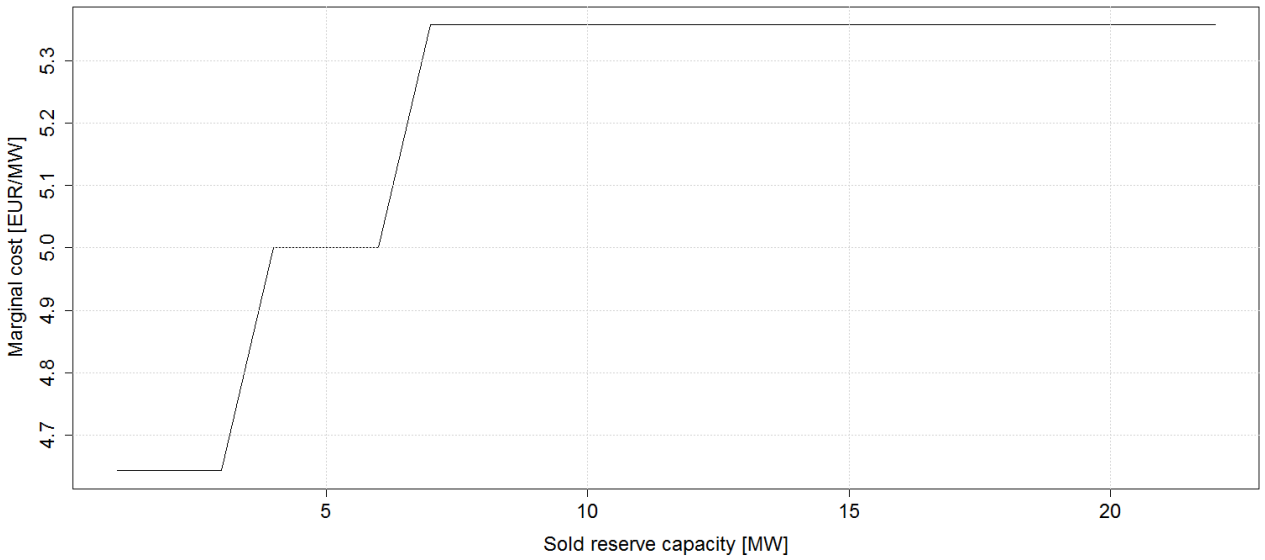

Figure 24 Marginal cost curve of capacity sales in week 45.

In Figure 25 the marginal cost associated with sales of different fixed levels of reserve capacity along the first simulated year is presented. We represent three different capacity levels; 1, 11 and 22 MW. The solid-drawn line represents capacity of 11 MW. As expected, we see that the marginal cost increases with increasing level of commitment. It is also evident that the marginal cost is much higher in certain periods of the year, which is related to low and high reservoir levels.

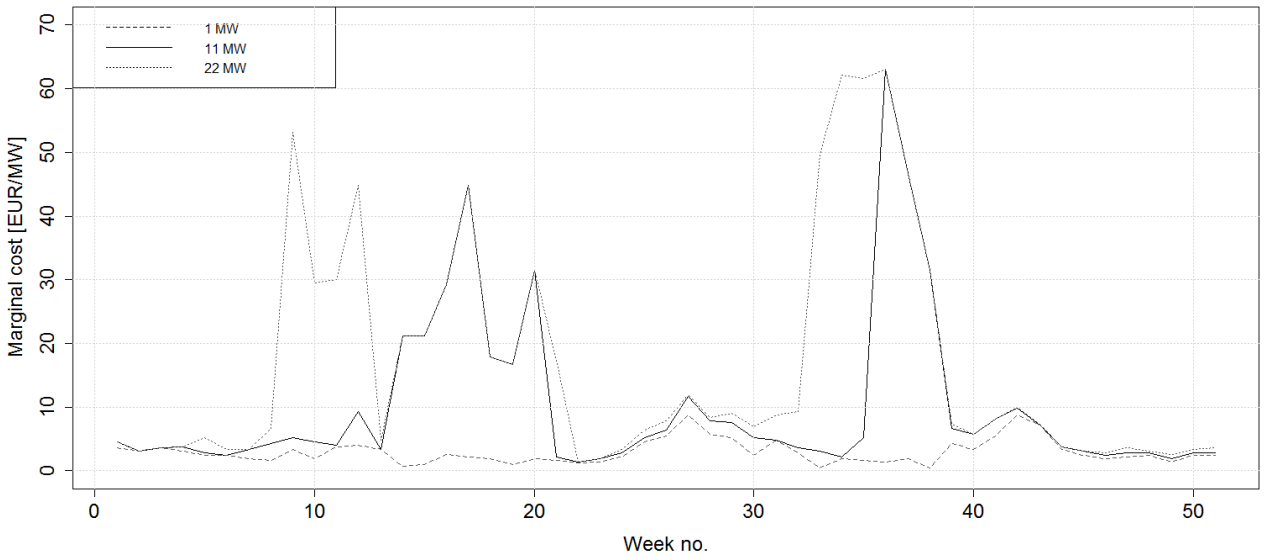

Figure 25 Marginal cost of delivering different amounts of reserve capacity in reserve block 3 (20:00-24:00) for the first year. The stapled lines correspond to 1 and 22 MW, and the solid-drawn line corresponds to 11 MW.

We wanted to test the dependence of the marginal cost shown in Figure 25 on the reserve capacity price level. For that reason we ran two additional cases, one where the capacity price profile was increased with a factor of 10, and one where it was decreased with a factor of 10. The expected marginal cost related to sales of 1 MW for the three cases (normal, high and low reserve capacity price) is shown in Figure 26. The radical variations in capacity price levels impacts the marginal cost curves, but not severely. Recall that the marginal cost does not explicitly depend on the reserve capacity price in a given decision stage. It will however implicitly depend on the reserve capacity price-level through the cuts/water values.

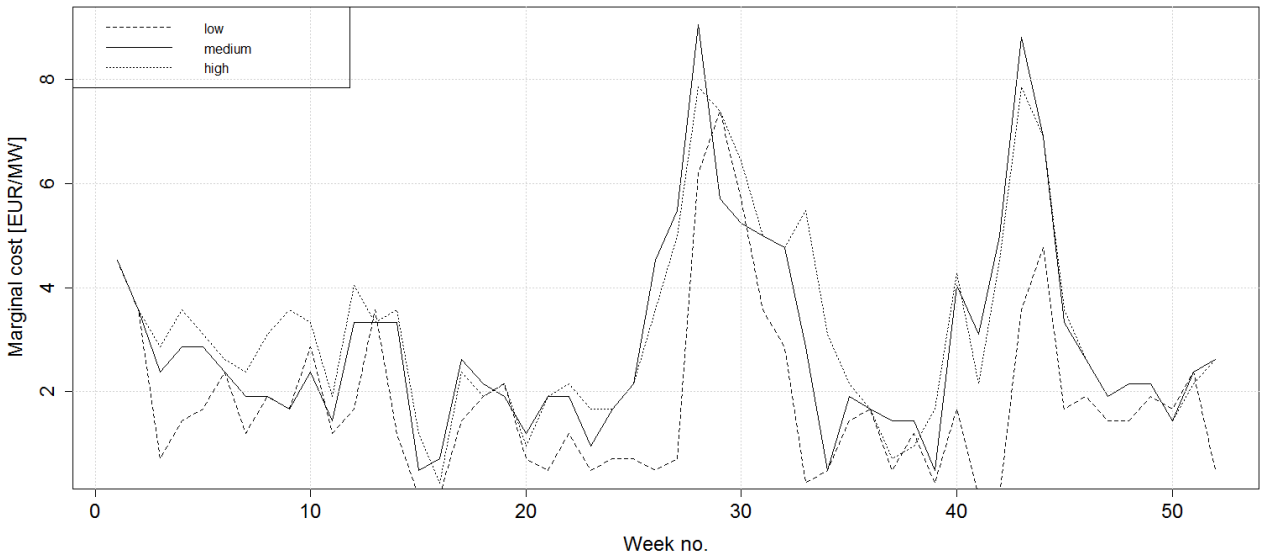

Figure 26 Marginal cost of delivering 1 MW reserve capacity in reserve block 3 (20:00-24:00) for the first year for different levels of reserve capacity prices.

The marginal costs in Figure 25 and Figure 26 fluctuate significantly between weeks. This can be explained by the fact that we have embedded a two-stage stochastic optimization problem within the SDDP/SDDP algorithm. That is, the marginal cost of the reserve capacity sold in the current week is found by evaluating the resulting commitment in the week ahead. The marginal cost therefore depends heavily on the energy price in the week ahead. The energy price fluctuations were displayed in Figure 8. We believe that the fluctuations will be less pronounced for a larger system and when considering more price nodes.

For comparison we ran two tests changing only the energy price input. In the first test we used two fixed energy price nodes throughout the period of analyses, one at 45 Euro/MWh and one at 55 Euro/MWh. Note that this is different from the original price input shown in Figure 8. The probabilities where set to 0.8 for a transition to the same price-level and 0.2 to a different level. The resulting marginal costs for reserve capacity are shown in Figure 27. As expected, the marginal costs vary less from week to week. In a second test we assumed a constant energy price of 50 Euro/MWh, see Figure 28. Sales of 1 MW reserve capacity comes "for free" most of the time. At the point of selling reserve capacity the model knows the energy price.

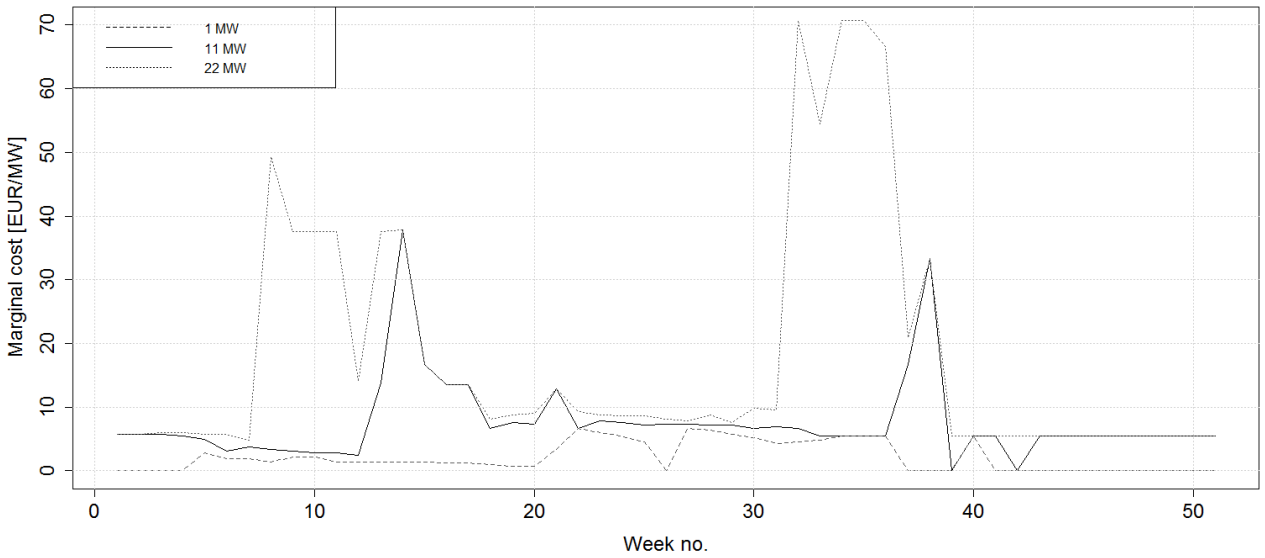

Figure 27 Marginal cost of delivering different amounts of reserve capacity in reserve block 3 (20:00-24:00) for the first year. Two energy price levels: 55 and 45 Euro/MWh.

As expected, the marginal cost curves in Figure 27 do not vary as much as in the original case (shown in Figure 25) between weeks.

In the second test we used a deterministic price, giving the results in Figure 28. The marginal cost of selling a small amount (1 MW) of reserve capacity is zero. In this case the energy price for the week ahead is perfectly known, but inflow is still uncertain when selling capacity. When comparing the marginal cost curves in Figure 28 with those in Figure 25 and Figure 27, we see that reduced uncertainty in the planning leads to lower expected marginal costs of reserve capacity. If we considered both the energy price and inflow as deterministic, we would allow the model to perfectly co-optimize the sales of energy and capacity. This will enable the model to perfectly plan when to deliver reserve capacity at zero cost.

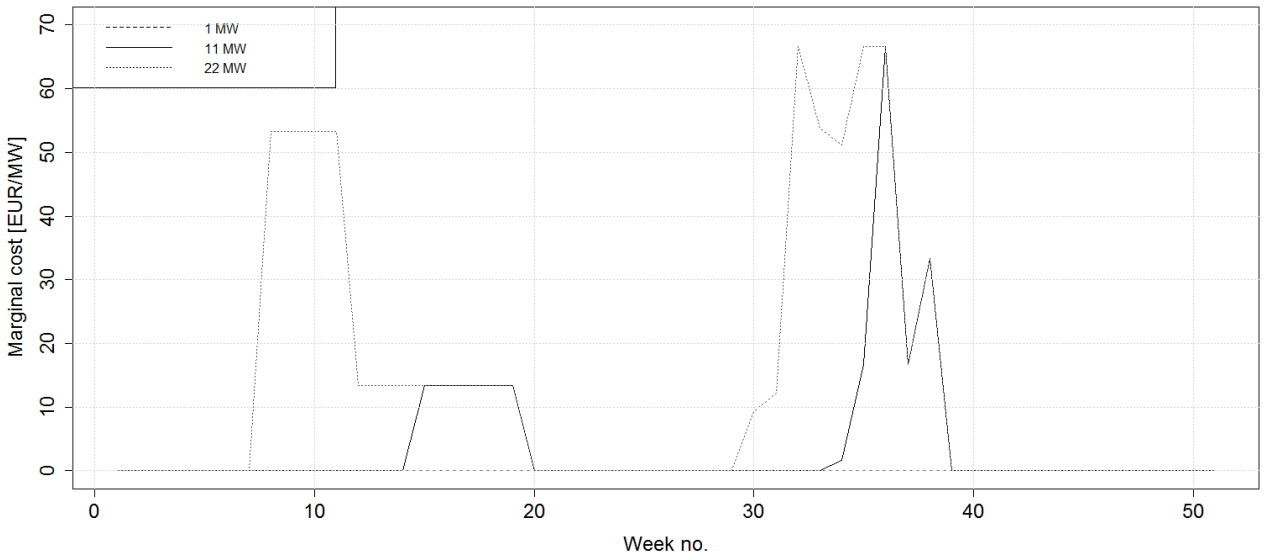

Figure 28 Marginal cost of delivering different amounts of reserve capacity in reserve block 3 (20:00-24:00) for the first year. Deterministic energy price: 50 Euro/MWh.

5 Conclusions

We have shown that sales of reserve capacity can be included in the combined SDP/SDDP method that is used in the program ProdRisk. The presented method assumes that reserve capacity is sold for the week-ahead and before knowing the energy price and inflow for that week. Alternatively, one could choose to co-optimize sales of energy and reserve capacity, but we think that the former choice better reflects the current market design.

When considering both energy and reserve capacity sales, water values change. This was elaborated analytically in section 3 and numerically in section 4.

In section 3, we showed that when the requirement for available down-regulation capacity is binding, it contributes to an increased water value. Conversely, we showed that when the requirement for available up-regulation capacity is binding, it contributes to a reduced water value. These findings where verified in section 4. From the case studies we saw that in cases with low reservoir level and high reserve capacity prices, the down-regulation capacity constraint becomes binding, contributing to an increase in water value (compared to the energy-only case). Conversely, with a high reservoir level and high reserve capacity prices, the up-regulation capacity constraint is binding, contributing to a decrease in water value (compared to the energy-only case). The same pattern was also found when increasing/decreasing the utilization time of the system.

The model needs reserve capacity prices, either represented as a deterministic or stochastic variable. The model user should also provide an estimate on the volumes that can be sold at the specified prices. Expected marginal costs for reserve capacity commitments seems to be an interesting by-product that can be extracted from the model. It reflects the lost revenue in the energy market when selling reserve capacity. It was shown that these costs heavily depends on the system state (reservoir level), but are relatively insensitive to the provided reserve capacity prices. For this reason, we argue that the presented functionality can be implemented in ProdRisk taking a dummy reserve capacity price and providing relatively stable estimates on the marginal cost of reserve capacity.

6 Possible implementation in ProdRisk

The presented method can be embedded in today's version of ProdRisk with modest programming effort. The following changes to the existing computer code are needed so that the program can:

- Read input regarding maximum reserve capacity per station or station group
- Include reserve constraints in the LP model
- Compute dual values for these constraints
- Include capacity sales in the cuts
- Extract results

We expect that the computation time will increase modestly since more constraints are included in the model. The convergence characteristics may also be slightly different since more state variables are introduced.

7 References

[1] A. Gjelsvik. Løysingsmetodikk i SDDP algoritmen. Technical Report TR F5560, SINTEF Energy Research, 2001.
[2] A. Gjelsvik, M. M. Belsnes, and A. Haugstad. An algorithm for stochastic medium-term hydrothermal scheduling under spot price uncertainty. In *Proc. 13th Power System Computation Conference*, Trondheim, Norway, 1999.
[3] A. Gjelsvik, B. Mo, and A. Haugstad. *Handbook of Power Systems I*, chapter Long- and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming, pages 33–55. Springer, 2010.
[4] A. Gjelsvik and S. W. Wallace. Methods for stochastic medium-term scheduling in hydro-dominated power systems. Technical Report TR A4438, SINTEF Energy Research, 1996.
[5] A. Helseth, M. Fodstad, and A. L. Henden. Balancing Markets and their Impact on Hydropower Scheduling – Review of Nordic Market Structures and Relevant Scheduling Methods. Technical Report TR A7558, SINTEF Energy Research, 2016.
[6] A. Helseth, M. Fodstad, and B. Mo. Optimal medium-term hydropower scheduling considering energy and reserve capacity markets. *IEEE Transactions on Sustainable Energy*, 2016.
[7] A. Helseth, B. Mo, M. Fodstad, and M. Hjelmeland. Co-optimizing sales of energy and capacity in a hydropower scheduling model. In *Proc. of IEEE PowerTech*, Eindhoven, The Netherlands, 2015.
[8] O. Wolfgang, A. Haugstad, B. Mo, A. Gjelsvik, I. Wangensteen, and G. Doorman. Hydro reservoir handling in Norway before and after deregulation. *Energy*, 34(10):1642–1651, 2009.

A Nomenclature

Index sets		**Stochastic variables**	
H	Set of hydropower reservoirs/stations	I_h	Sum weekly inflow to reservoir h
S_h	Set of discharge segments for station h	$\lambda^C_{b,t}$	Weekly average reserve capacity price
Ω_h	Set of reservoirs upstream reservoir h	λ^E_p	Weekly average energy price
K	Set of time steps within the week		
B	Set of reservation blocks within the week		
Parameters		**Decision variables**	
π_{phl}	Coefficient for reservoir level for price node p, reservoir h and cut l	e_{kh}	Generated electricity in time step k for station h
μ_{pbl}	Coefficient for capacity sales for price node p, block b and cut l	w_h	Artificial water supply to reservoir h
β_{pl}	Right-hand side for price node p and cut l	$\alpha_{p,t+1}$	Future expected profit for price node p and week $t + 1$
$P^{\max}_h, P^{\min}_h$	Max./Min. capacity in station h	v_{kh}	Volume in time step k for reservoir h
$R^{\max}_h$	Max. capacity reservation for station h	q^D_{khs}	Discharge in time step k through station h at segment s
γ_h	Factor limiting the use of spinning reserves	q^S_{kh}	Spillage in time step k from reservoir h
τ_b	Total duration of reservation block b	q^B_{kh}	Bypass in time step k from reservoir h
$\varsigma^C_{b,t}$	Capacity price scaling coefficient for block b	r_{kh}	Allocated capacity in time step k for station h
ς^E_k	Energy price scaling coefficient for time step k	$c_{b,t+1}$	Sold capacity for block b in week $t + 1$
φ	Cost of artificial water		
τ_k	Duration of time step k		
η_{hs}	Energy equivalent for station h and discharge segment s		
$V^{\max}_{kh}, V^{\min}_{kh}$	Max./Min. limit for reservoir h		

Technology for a better society

www.sintef.no