

Revisiting hierarchical planning for hydropower plant upgrades using semi-analytical policies and reinforcement learning

Andreas Kleiven¹, Selvaprabu Nadarajah,² and Stein-Erik Fleten¹

¹Department of Industrial Economics and Technology Management,
Norwegian University of Science and Technology, Trondheim, Norway

²College of Business Administration,
University of Illinois at Chicago, Chicago, Illinois 60607

Abstract

Hydropower plants provide flexibility and storage to support the penetration of renewable energy sources needed to meet climate goals. Investments to upgrade their capacity depend on accurate valuation. Such strategic valuation in principle depends on long-term market price movements, tactical capacity allocation, and capacity bids that respond to short-term price fluctuations. Given the complexity of this holistic problem, hierarchical planning is commonplace, where investment models simplify tactical capacity allocation decisions and ignore the value of short-term production flexibility. We formulate a novel investment model that accounts for these aspects. While our problem is complex, we show how a combination of price modeling, informed by empirical analysis, and the use of reinforcement learning to solve for capacity allocation can lead to insightful semi-analytical investment policies. In particular, these policies highlight that capacity investment is supported at lower power prices when the short-term variability of these prices increases, that is, when the value of short-term production flexibility is higher. A numerical study based on real operational and market data shows that valuations from our model can be computed efficiently. Our findings suggest that investment models enabled by reinforcement learning that value the operational flexibility of production assets at long and short time scales can significantly help promote additional capacity in hydro power. The tools we develop are potentially relevant for analogous valuation of investments in other renewable energy production assets.

1. Introduction

Hydropower plants with storage reservoirs are flexible and can help to integrate the increasing amount of wind and solar. However, many hydropower plants rely on old technology from the large-scale hydropower projects in the mid-20th century, which were designed to operate under different market conditions (IRENA 2015, EIA 2017). In addition to an aging hydropower fleet, most of the economically viable hydropower potential in developed regions, such as Europe, Canada, and the United States, is already exploited. Thus, upgrading existing hydropower plants is becoming increasingly relevant.

Investment in hydropower capacity is costly, and to support these decisions, accurate mathematical models and methods for the calculation of operational revenues associated with investment opportunities are crucial. Producers typically aim at maximizing the estimated market value of their assets (Wallace and Fleten 2003, Nadarajah and Secomandi 2021). This involves the estimation of the real options value of dynamically taking profit-maximizing decisions in response to exogenous information (Dixit and Pindyck 1994, Duffie 2010), which means establishing production

and investment policies in the presence of uncertain factors such as weather and risk-adjusted prices (Dimoski et al. 2019). In addition, hydropower operators can utilize hourly variations in prices operationally. Uncertainty in weather and prices, and the short time scale required to capture the inherent flexibility that the producer has to respond to short-term price fluctuations, make capacity valuation a complex problem.

Hierarchical planning provides a framework for decision making at different organizational levels. Anthony (1965) classifies three categories: Strategic planning, tactical planning, and operations control. At lower levels, detailed decisions determine day-to-day operations. At higher levels, aggregate decisions determine resource utilization and expansion projects. This decoupling breaks down complex and coupled problems into tractable optimization formulations that can be solved for each level. Hierarchical planning has been widely used in applications such as manufacturing and distribution of products, vehicle routing and scheduling, and production scheduling, among other applications (Dempster et al. 1981). A notable characteristic of planning in the setting of renewable energy production scheduling and capacity installment is the interaction with financial markets and weather. Thus, the valuation of capacity is intimately tied to market and weather dynamics.

In this work, we formulate the capacity upgrade problem as a Markov decision process (MDP). The MDP accounts for detailed production decisions for the valuation of resources in the presence of market and weather dynamics. Because of the high-dimensional state space, long time horizon, non-convex action set, and non-linear cost structure, the MDP is computationally intractable. Based on the structural properties of the MDP, we propose to solve the problem using reinforcement learning combined with real options analysis to obtain insightful semi-analytical policies. Our analysis shows that the integration of tactical planning and operations control significantly enhances the resource value and thus affects long-term strategic planning policies. As a result, higher capacity investments should be undertaken sooner if the short-term flexibility to respond to price fluctuations is accounted for. We further analyse a situation where hourly prices vary more, and we find that assumptions regarding hourly price variations have a substantial impact on capacity investment policies.

1.1 Novelty and related work

We contribute to the growing literature on decision making in renewable energy from a price-taker perspective (Fleten et al. 2007, Boomsma et al. 2012, Liu et al. 2019). Specifically, we study capacity investments in existing hydropower plants. Related work include Bøckman et al. (2008) who analyze investments in new small hydropower projects under long-term price uncertainty, assuming the resource value is exogenously given, and Andersson et al. (2014) who combine production scheduling and investment valuation. Compared to these works, our production scheduling and resource valuation are considerably more sophisticated. In particular, we include the short-term operational control aspect for better capturing the short-term flexibility. Furthermore, we provide bounds on the value of the optimal production schedules, and we obtain semi-analytical investment policies for hydropower plant upgrades, which allows us to analyze the effect of short-term price variations on long-term capacity investments.

Our second contribution is to the literature on hierarchical planning (Anthony 1965, Bitran and Tirupati 1993, Dempster et al. 1981, Lenstra et al. 1984). More specifically, we contribute to the literature on hydropower reservoir management and production planning (Gjelsvik et al.

2010, Shapiro et al. 2013, Löhndorf et al. 2013). Because of computational challenges, the reservoir management problem is solved at different time scales, where the output of a seasonal problem, i.e. tactical decisions, are used to estimate the end of horizon value of the short-term problem, i.e. operations control (Wallace and Fleten 2003). We propose to formulate the combined short-term and seasonal planning problem as an MDP for resource valuation, integrating tactical planning and operations control. This accounts for both uncertainty in weekly prices and inflows, and the flexibility to quickly ramp up and down production, which we show are both important features when considering hydropower capacity investments.

Our third contribution is to the reoptimization literature for obtaining feasible and near-optimal policies in multistage stochastic problems. The reoptimization heuristic has been extensively applied by energy storage practitioners and academic researchers (Lai et al. 2010, Wu et al. 2012, Nadarajah and Secomandi 2018, Löhndorf and Wozabal 2021), and information relaxations and duality theory is a popular approach for bounding the policy value (Brown et al. 2010). Lai et al. (2010) find that the rolling intrinsic policy is near-optimal in gas storage applications, and Secomandi (2015) provides theoretical support for this. However, in the application of hydropower, there is exogenous factors that cannot be perfectly hedged, i.e. inflow, which may lead to poorer performance of the reoptimization heuristic. The state-of-the-art algorithm for solving the seasonal hydropower problem is stochastic dual dynamic programming (Pereira and Pinto 1991), which has convergence guarantees under certain assumptions (Philpott and Guan 2008). However, in our setting, the reoptimization heuristic is suitable for tackling the integrated tactical planning and operations control MDP. We analyse numerically situations where the reoptimization heuristic provides, or fails to provide, a near-optimal policy.

Our fourth contribution is to the real options literature on irreversible timing options (Dixit and Pindyck 1994). We consider a one-time investment opportunity in production capacity, where the capacity level is optimized upon commencement (Dangl 1999, Fleten et al. 2007, Hagspiel et al. 2016, Huisman and Kort 2015). We study an existing hydropower plant, which means that additional cashflows from operations start being generated after the investment. Often, the relationship between capacity and production output is often subject to coarse approximations in the real options literature on irreversible timing options. We show how a combination of price modeling and the use of reinforcement learning for operations planning can lead to insightful semi-analytical capacity investment policies.

1.2 Paper structure

The paper is structured as follows: In Section 2 we present the common industry planning approach from the perspective of hydropower producers. In section 3 we present a capacity upgrade MDP. The solution approach is presented in Section 4. Section 5 contains numerical results from case studies. Finally, concluding remarks are provided in Section 6.

2. Industry practice through a hierarchical-planning lens

In this section, we overview the planning approach in the hydropower industry, illustrated in Figure 1, based on our interactions with companies in this sector. While each company has its own internal processes, which may vary, our goal is to focus on the core ideas and highlight the heightened importance of valuation in this context.

Hydro power capacity upgrade decisions often involve net present value (NPV) analyses over

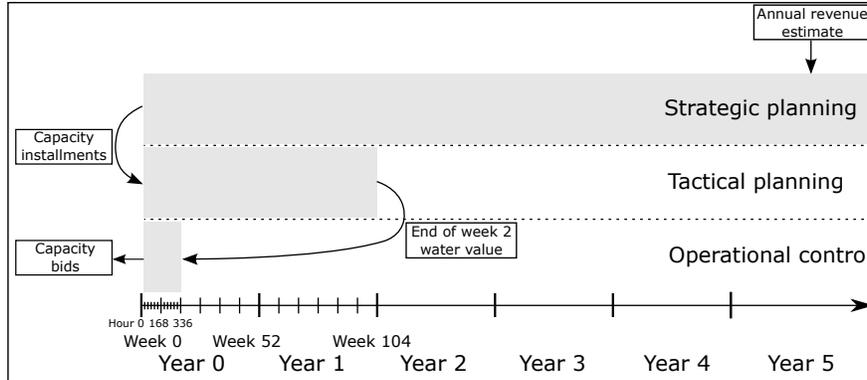


Figure 1: Common industry practice for hydropower planning.

a horizon spanning about thirty years. This is illustrated in the upper level in Figure 1. In particular, NPV calculations are performed for different capacity levels. The annual cash flow estimates needed for these calculations correspond to an extrapolation of revenues either from a representative week obtained by heuristically using short-term operational models that account for short-term flexibility, or from a representative year using medium-term tactical planning models that account for price and weather uncertainty. Long-term trends in prices and inflow of water are considered by repeating the NPV analysis for several long-term market and climate scenarios. In each case, the NPV is compared against an internal investment hurdle to determine if the capacity should be upgraded and by what amount. The outcome of this strategic planning exercise is the upgraded capacity of the plant, which is a critical input to the tactical planning phase.

The goal of tactical planning is to obtain an accurate estimate of the marginal (production) value of water in the short-term, for instance, a week ahead. This is illustrated in the middle level in Figure 1. This value is the expected price at which power generated from a unit of water can be sold for in the market if the timing of this sale is optimized accounting for reservoir and production capacities as well as inflow and price dynamics. The water value is expressed in monetary units per MWh (e.g., \$ per MWh or € per MWh). To get an accurate estimate of the marginal value of water, the tactical production decisions need to be planned over a 1-3 year horizon. Owners of large reservoirs relative to annual expected inflows need a long time horizon to avoid end-of-horizon effects on the near-term marginal value of water, while small reservoirs that quickly get emptied need a shorter time horizon. The tactical planning phase accounts for uncertainty in prices and inflows, partly at the expense of physical system details for computational tractability. Stochastic models for uncertain prices are either based on the financial market or fundamental modelling, or a combination. For inflows, producers may use up-to-date probabilistic weather forecasts for the near future, before transitioning to a stochastic model that includes time of the year as a feature and intertemporal effects. Given these models, tactical production decisions can be optimized, from which the marginal water value at any point in time and state can be derived. This serves as the end of horizon value for the short-term models. Therefore, week-ahead water values are typically reestimated every week using updated forecasts for prices and weather. For more details on tactical planning in hydropower, we refer to [Gjelsvik et al. \(2010\)](#) and [Pérez-Díaz et al. \(2020\)](#).

The short-term model aims at providing decision support for day-to-day operations. This is illustrated in the lower level in Figure 1. At this time scale, the decision problem is often

considered as deterministic, which allows for a more precise description of the physical system. Still, stochasticity in prices and inflows at a longer time scale enters through the estimated state-dependent water values from the tactical planning level. The output from the short-term model is used for constructing bids in the day-ahead market. Every day, individual producers submit bids to the market operator in the form of price-volume pairs for each hour of the coming day. Therefore, the short-term models are rerun daily based on updated information on weather and markets for efficient use of water. This daily reoptimization means that the aggregated short-term operational decisions do not necessarily correspond to the associated tactical decisions at the higher level. For more details on short-term operational modeling in hydropower, see, e.g. [Conejo et al. \(2002\)](#) or [Belsnes et al. \(2016\)](#).

The planning flow summarized above is an example of hierarchical planning in renewable energy production and has several advantages that are well known in the standard use of such planning in e.g. manufacturing and distribution of products, or job shop design and scheduling ([Dempster et al. 1981](#)). It breaks down an otherwise coupled and challenging planning problem into smaller manageable strategic, tactical, and operational sub problems, which has several advantages. First, the decisions in each subproblem of the hierarchy can be modeled more accurately while still leading to solvable optimization formulations. Second, the hierarchical approach facilitates planning at different time scales. For instance, the strategic problem, which often is a one-time decision that requires a planning horizon of several decades, determines the capacity upgrade amount and only approximately captures the effect of day-to-day operations and tactical planning decisions through the cash flow estimate of a representative week or year. Thus, the detailed planning of production decisions at lower levels, which needs to be implemented at some future point in time, can be postponed. These decisions can then benefit from the unfolding of information and up-to-date forecasts regarding the market and inflow, which are unavailable at the time the long-term strategic planning model is solved.

The hierarchical structure used in practice, however, can be improved from a valuation standpoint, while retaining the flexibility it provides for operational decision making. This is easiest to explain by focusing on investment planning. We focus on two main aspects which can be improved from a valuation standpoint: First, the integration of approximate operations control and medium-term capacity allocation for resource valuation, and second, the real options value associated with the timing flexibility of capacity installments. As explained above, the water value estimates obtained from solving the seasonal resource valuation problem are linked to capacity constraints. However, if the operational flexibility on a shorter time scale is ignored during the water value estimation, the value of possessing capacity may be underestimated. On the other hand, the short-term problem needs an estimate of the marginal end-of-horizon water value. This makes the integration of the short-term production model and seasonal resource valuation model relevant from a valuation standpoint. Furthermore, there is managerial flexibility in the timing of capacity installments. Thus, there may be an opportunity cost of installing an investment project if acting according to the NPV and internal investment hurdle rule. This aspect has been addressed in the academic literature by [Bøckman et al. \(2008\)](#). However, they consider pure strategic planning of new hydropower plants, assuming annual cashflows from a set of capacity alternatives are exogenously given. We focus on the integration of the hierarchical levels for valuation of additional capacity installments in existing hydropower plants. Finally, although we focus on the investment

problem, the same idea also helps improve the water valuation by the proposed integration of the short-term model and the seasonal planning model.

3. Hydropower capacity upgrade as an MDP

In this section, we present the hydropower capacity upgrade MDP. We first present the dynamics of exogenous factors, before we describe the long-term optimal stopping and upgrade problem. Next, we focus on the resource valuation, which determines the cashflows associated with capacity upgrades. We propose two versions for the resource valuation; with and without the operations control aspect, i.e. short-term within-week flexibility. We formulate the problem in energy equivalents, which means that capacities and volumes are already transformed into units of MWh ([Arvanitidits and Rosing 1970](#)).

3.1 Data and stochastic model

Exogenous factors that affect hydropower planning are prices and inflows to the reservoir. We denote the price and inflow at time $t \in [0, \infty)$ by S_t and Z_t , respectively. We study a price-taking hydropower operator that aims at maximizing the market value of the assets. Therefore, we calibrate a stochastic model using futures prices on electricity. The price of these contracts is based on the market expectations of the average price of electricity over some period of time. Thus, the first step is to construct smooth synthetic futures curves. We do this by applying the method by ([Benth et al. 2007](#)). The futures price data set is obtained from [Montel \(2021\)](#) and consists of futures contracts traded in the period from 2012 to 2018 with maturities between 1 month and 5 years. Given historical synthetic weekly futures prices, we calibrate the two-factor stochastic model

$$\ln S_t = \phi_1 \cos\left(\frac{2\pi t}{\theta}\right) + \phi_2 \sin\left(\frac{2\pi t}{\theta}\right) + \chi_t + \xi_t, \quad (1)$$

$$d\chi_t = (-\kappa_\chi \chi_t - \lambda_\chi)dt + \sigma_\chi dz_\chi \quad (2)$$

$$d\xi_t = \left(\mu_\xi - \lambda_\xi - \frac{1}{2}\sigma_\xi^2\right)dt + \sigma_\xi dz_\xi, \quad (3)$$

where $dz_\chi dz_\xi = \rho_{\chi\xi} dt$. The trigonometric functions account for seasonality, where θ is the number of periods in one cycle, in our case 52 weeks. The stochastic factors account for long-term price movements through a geometric Brownian motion and short-term price deviation through a mean-reverting process ([Schwartz and Smith 2000](#)). The parameters ϕ_1 and ϕ_2 determine seasonal effects, and κ_χ accounts for mean reversion in prices. Parameters λ_χ and λ_ξ are risk premiums, and σ_χ and σ_ξ are volatility in short-term prices and long-term prices, respectively. The drift of long-term log prices are given by $\mu_\xi - \lambda_\xi - \frac{1}{2}\sigma_\xi^2$, and the correlation between short-term and long-term prices is $\rho_{\chi\xi}$.

As described in the previous section, practitioners simulate prices either based on fundamental models or market models. For the latter, the two-factor model that we apply is suitable. Monte Carlo simulations can be obtained from the scheme

$$S_t = \exp\left(\phi_1 \cos\left(\frac{2\pi t}{\theta}\right) + \phi_2 \sin\left(\frac{2\pi t}{\theta}\right) + \chi_t + \xi_t\right) \quad (4)$$

$$\xi_{t+\Delta t} = \xi_t + \left(\mu_\xi + \lambda_\xi - \frac{1}{2}\sigma_\xi^2\right)dt + \sigma_\xi \sqrt{\Delta t} \epsilon_1 \quad (5)$$

$$\chi_{t+\Delta t} = \chi_t e^{-\kappa_\chi \Delta t} - \frac{\lambda_\chi}{\kappa_\chi} (1 - e^{-\kappa_\chi \Delta t}) + \sigma_\chi \sqrt{\frac{1 - e^{-2\kappa_\chi \Delta t}}{2\kappa_\chi}} \epsilon_2, \quad (6)$$

where ϵ_1 is standard Gaussian and ϵ_2 is Gaussian with mean $\rho_{\chi\xi}\epsilon_1$ and variance $1 - \rho_{\chi\xi}^2$. By letting $\Delta t = \frac{1}{52}$ years, weekly prices can be simulated. For hourly prices, we need to estimate time-of-week effects for the short-term operations control aspect. This time scale is too short to be observed in the financial market. We therefore estimate the hourly price profile using historical spot price data from NordPool (2021), which is the market operator in the Nordic countries. We use the spot prices to estimate the expected deviation from the mean in each hour of the week. Typically, the price during the day is above the weekly mean and prices during the night are below the weekly mean. We denote the within-week price profile at time t by the vector α_t . Thus, the price in each hour of week t , given a simulated weekly log price mean $\ln S_t$, is given by $\ln S_t + \alpha_t$. An example of such a profile is provided in Section 3.4 where we illustrate the operations control aspect.

The inflow data set is supplied by a hydropower producer located in the western part of Norway. The data set consists of weekly inflow observations in the period from 2009 to 2018. We use a similar model as in Gjelsvik et al. (2010), which also is common among practitioners,

$$Z_t = \max\{(\bar{\mu}_t + \bar{\sigma}_t \zeta_t), 0\} \quad (7)$$

$$d\zeta_t = -\kappa_\zeta \zeta_t dt + \sigma_\zeta dz_\zeta, \quad (8)$$

where $\bar{\mu}_t$ is an estimated of the mean at time t , $\bar{\sigma}_t$ an estimate of the standard deviation, and σ_ζ determines the volatility of inflow deviations from the mean. The reason for the normalization is because inflows typically possess strong seasonal characteristics, both in mean and variance. The stochastic factor in (8) accounts for inter-temporal effects through mean-reversion parameter κ_ζ and can be simulated as follows,

$$\zeta_{t+\Delta t} = \zeta_t e^{-\kappa_\zeta \Delta t} + \sigma_\zeta \sqrt{\frac{1 - e^{-2\kappa_\zeta \Delta t}}{2\kappa_\zeta}} \epsilon_3, \quad (9)$$

where ϵ_3 is a standard Gaussian variable. As opposed to prices, there are no apparent time-of-week effects for inflows. Therefore, we uniformly allocate the accumulated weekly inflow to each hour of the week, which we elaborate on when we present the operations control aspect.

Throughout this section, we use Greek letters for stochastic variables and parameters of the stochastic processes (parameters in (1)-(3), (7)-(8), and the price profile α_t). We use bold letters for vectors and denote the exogenous state vector at time t by $\omega_t = (\chi_t, \xi_t, \zeta_t) \in \mathbb{R}_+^3$. Furthermore, we use lower case letters for decision variables and endogenous states and capital letters for derived exogenous states (S_t and Z_t) and for parameters of the decision model, which we introduce next.

3.2 Optimal stopping and upgrade decision

At any time $t \in [0, \infty)$, additional production capacity can be installed. The plant has an initial production capacity q_0 and reservoir capacity R . We denote the one-time unknown upgrade time by τ . The upgrade decision at time τ is a scalar denoted $u \in (0, R - q_0]$, where R is the maximum reservoir capacity. Thus, the endogenous strategic state is

$$q_t(\tau, u) = \begin{cases} q_0 & \text{if } t < \tau, \\ q_0 + u & \text{if } t \geq \tau. \end{cases} \quad (10)$$

Choosing $u > 0$ entails an immediate cost

$$K(u) = Ae^{Bu}, \quad (11)$$

where the parameter A determines the cost of initiating capacity upgrade, and the parameter B determines the relationship between cost and additional capacity. This cost function has been used in the academic literature (Bøckman et al. 2008). In practice, it is difficult to specify and estimate a cost function for capacity installments. However, some general features applies, such as an increasing marginal cost of installing capacity. The goal is to find a strategic planning policy, i.e. timing τ and capacity installment u . The optimal stopping problem with capacity upgrade can be formulated as

$$H(q_0, \boldsymbol{\omega}_0) = \max_{\tau \in [0, \infty), u \in (R - q_0]} \mathbb{E} \left[\int_0^\infty \gamma_1^t V_t(\boldsymbol{\omega}_t, q_t(\tau, u)) dt - \gamma_1^\tau K(u) \middle| \boldsymbol{\omega}_0 \right], \quad (12)$$

where $\gamma_1 = e^{-r}$ is the continuously compounded discount rate and r is the risk-free rate. The capacity function $q_t(\tau, u)$ is defined in (10) and the cost function $K(u)$ is defined in (11). The decision when to upgrade and by how much requires an annual valuation of the resource (water) as a function of time and capacity. This is denoted $V_t(\boldsymbol{\omega}_t, q_t)$ in (12). To obtain this valuation, we solve a production scheduling problem over the interval $[t, t + \Delta t]$, where Δt is an annual interval. We explain this next.

3.3 Resource valuation and production decisions

We obtain the resource valuation $V_t(\boldsymbol{\omega}_t, q_t(\tau, u))$ in (12) at a given time t and a given capacity installment at time t by solving a production scheduling problem for several capacity alternatives. Thus, capacity is a parameter in the resource valuation phase. We therefore introduce $Q := q_t(\tau, u)$, while keeping t fixed throughout this section. We first explain the resource valuation with weekly decision periods without short-term flexibility.

Resource valuation without short-term flexibility: The resource valuation corresponds to the tactical planning level in Section 2 and illustrated in Figure 1. This is often referred to as *medium-term hydropower scheduling* in the academic literature (Flatabø et al. 1998, Gjelsvik et al. 2010). Decisions are made weekly, which we denote by index $i \in \mathcal{I}_t = [0, 1, \dots, nI]$, where $I := \Delta t$ is a one-year horizon. To avoid end of horizon effects, the tactical planning horizon for resource valuation is a multiple of n years. The endogenous tactical state in week i is the reservoir volume $l_{t,i} \in [0, R]$, and the decision is weekly production $x_{t,i} \in \mathcal{X}_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i})$. Executing $x_{t,i}$ in stage i entails an immediate reward

$$r_{t,i}^W(x_{t,i}, \boldsymbol{\omega}_{t,i}) = S_{t,i} x_{t,i}, \quad (13)$$

which is the product of price and production volume, i.e. revenue from electricity sales. Superscript W abbreviates *week* and refers to the length of the decision period. The price $S_{t,i}$ is defined in (1).

When water is used for electricity generation, the reservoir volume gets updated as follows,

$$l_{t,i+1} = f(l_{t,i}, x_{t,i}, \boldsymbol{\omega}_{t,i}) = l_{t,i} - x_{t,i} + Z_{t,i} - z_{t,i} \quad i \in \mathcal{I}_t, \quad (14)$$

where $z_{t,i}$ is spillage. Spillage happens if unexpectedly much inflow arrives at a given stage, which cannot be stored because of the limited reservoir and/or production capacity. This transition function ensures the energy balance in time. It states that the reservoir volume at stage i needs to be the sum of the reservoir volume in stage i , $l_{t,i}$, and incoming inflow $Z_{t,i}$ minus the sum of production $x_{t,i}$ and potential spillage $z_{t,i}$. The endogenous state update is illustrated in Figure 2. In the upper level, the capacity gets updated. In the lower level, the capacity is kept fixed during the resource valuation, which happens in the interval $[t, \Delta t]$. In the lower level, the reservoir volume gets updated for a given upper level time t and capacity $q_t = Q$.

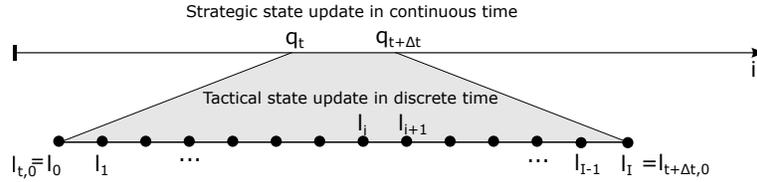


Figure 2: Illustration of the planning horizon. In the upper level, the capacity gets updated in continuous time $t \in [0, \infty)$. In the lower scale, the reservoir volume gets updated in discrete time $i \in \mathcal{I}_t$.

The generation capacity Q determines the boundary of the feasible stage- i action set which is given by

$$\mathcal{X}_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q) = \{ \quad l_{t,i} - x_{t,i} + Z_{t,i} - z_{t,i} \leq R \quad (15a)$$

$$x_{t,i} \leq l_{t,i} + Z_{t,i} \quad (15b)$$

$$x_{t,i} \leq Q \quad (15c)$$

$$x_{t,i}, z_{t,i} \geq 0 \}, \quad (15d)$$

where constraint (15a) ensures that the reservoir volume is less than the maximum reservoir capacity and constraints (15b)-(15d) are bounds on the maximum and minimum production and spillage. The resource value in stage i and state $(l_{t,i}, \boldsymbol{\omega}_{t,i})$ for capacity Q is given by a stochastic dynamic program (SDP). We denote by $V_{t,i}^W(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)$ and $W_{t,i}(\cdot, \boldsymbol{\omega}_{t,i}; Q)$ the value function and continuation function of the SDP, respectively, at the weekly time scale. The SDP can be formulated as

$$V_{t,i}^W(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q) = \max_{(x_i, z_i) \in \mathcal{X}_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)} r_{t,i}^W(x_{t,i}, \boldsymbol{\omega}_{t,i}) + \gamma_2 W_{t,i}(l_{t,i+1}, \boldsymbol{\omega}_{t,i}; Q) \quad (16)$$

$$W_{t,i}(\cdot, \boldsymbol{\omega}_{t,i}; Q) = \mathbb{E}(V_{t,i+1}^W(\cdot, \boldsymbol{\omega}_{t,i+1}; Q) | \boldsymbol{\omega}_{t,i}), \quad (17)$$

with $V^W(l_{t,nI}, \boldsymbol{\omega}_{t,nI}; Q) = 0$ for each $l_{t,nI}$, and where $\gamma_2 := e^{-r\Delta i}$ with $\Delta i = \frac{1}{52}$ and r is the risk-free rate. The immediate reward $r_{t,i}^W(x_{t,i}, \boldsymbol{\omega}_{t,i})$ is defined in (13). Next, we present the resource valuation with short-term flexibility to exploit within week variations in prices.

Resource valuation with short-term flexibility: Finally, we account for the short-term flexibility of the operations control aspect, often referred to as *short-term hydropower scheduling* in the academic literature (Flatabø et al. 1998, Borghetti et al. 2008, Belsnes et al. 2016). Our approach in this section is different from the procedure we described in Section 2 where water values from the tactical phase serve as end-of-horizon condition for the short-term operational problem. Instead, we integrate the short-term operational aspect in the tactical planning phase. As explained Section 2, producers submit price-volume pairs for every hour the next day to the market operator. We approximate this process by introducing hourly decision periods $j \in \mathcal{J} = \{0, 1, \dots, J\}$. We augment the stage i decision vector and define the hourly decision vector $\mathbf{y}_{t,i} = \{y_{t,i,j}, j \in \mathcal{J}\}$ for week i . Similarly, for spillage we define $\mathbf{v}_{t,i} = \{v_{t,i,j}, j \in \mathcal{J}\}$. Furthermore, we introduce variables $\mathbf{m}_{t,i} = \{m_{t,i,j}, j \in \mathcal{J}\}$ for keeping track of the reservoir volume between hours of the week. As explained in Section 3.1, the price for each hour j in week i and year t , given the price realization $S_{t,i}$ is obtained by a price profile $\boldsymbol{\alpha}_{t,i} = (\alpha_{t,i,j}, j \in \mathcal{J})$, which determines the price in every hour of the week. Thus, the log price in hour j is $\ln S_{t,i,j} = \ln S_{t,i} + \alpha_{t,i,j}$. For an illustration of $\boldsymbol{\alpha}_{t,i}$, see Figure 13 in Appendix B. The hourly inflow is obtained by uniform allocation of the (realized) weekly inflow, $Z_{t,i,j} = \frac{1}{J}Z_{t,i}$. The modified revenue in stage i becomes

$$r_{t,i}^H(\mathbf{y}_{t,i}, \boldsymbol{\omega}_{t,i}) = \sum_{j \in \mathcal{J}} \gamma_3^j y_{t,i,j} S_{t,i,j}, \quad (18)$$

where γ_3 discounts cashflows within a stage i , and superscript H abbreviates *hour*, which refers to the decision granularity. The endogenous state update becomes

$$l_{t,i+1} = f(l_{t,i}, \mathbf{y}_{t,i}, \boldsymbol{\omega}_{t,i}) = l_{t,i} - \sum_{j \in \mathcal{J}} y_{t,i,j} + \sum_{j \in \mathcal{J}} Z_{t,i,j} - \sum_{j \in \mathcal{J}} v_{t,i,j} \quad i \in \mathcal{I}_t, \quad (19)$$

which states that the reservoir volume in the next stage is the reservoir volume from the previous stage plus the incoming inflow in every hour of the week minus the sum of production in each hour of the week minus the potential spillage in each hour. The stage i action set in (15a)-(15d) is modified to

$$\mathcal{Y}_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q) = \{ \quad m_{t,i,0} = l_{t,i} \quad (20a)$$

$$m_{t,i,j+1} = m_{t,i,j} - y_{t,i,j} + Z_{t,i,j} - v_{t,i,j} \quad j \in \mathcal{J} \setminus \{0\} \quad (20b)$$

$$m_{t,i,j} - y_{t,i,j} + Z_{t,i,j} \leq R + v_{t,i,j} \quad j \in \mathcal{J} \quad (20c)$$

$$y_{t,i,j} \leq m_{t,i,j} + Z_{t,i,j} \quad j \in \mathcal{J} \quad (20d)$$

$$y_{t,i,j} \leq \frac{Q}{J} \quad j \in \mathcal{J} \quad (20e)$$

$$y_{t,i,j}, v_{t,i,j} \geq 0 \quad j \in \mathcal{J}. \quad (20f)$$

The reservoir volume transition in (20b) is illustrated in the lower level in Figure 3. We emphasize that there still are only two levels where decisions are made: The upper strategic level for capacity installments, and the lower tactical and operational level for production decisions which is used for resource valuation. The important parameter/variable that links the levels of the hierarchy is Q in (20e), which is a parameter in the lower resource valuation level, and a decision variable in

the upper strategic level. Our formulation is based on energy equivalents. Therefore, Q is measured in units MWh per length of the decision periods. We specify it as MWh per week, which is why capacity is allocated uniformly to hours of the week, as seen in the right hand side in (20e).

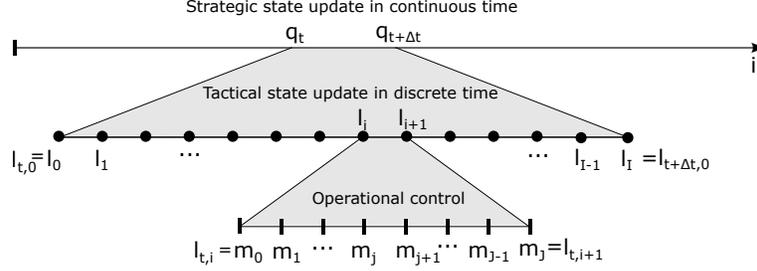


Figure 3: Illustration of the planning hierarchy and time scales. In the upper level, the capacity gets updated in continuous time $t \in [0, \infty)$. In the middle scale, the reservoir volume gets updated in discrete time $i \in \mathcal{I}_t$. The lowest scale illustrates the operations control aspect in discrete time $j \in \mathcal{J}$. Subscripts for fixed indices are omitted in the figure.

We denote by $V_{t,i}^H(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)$ and $C_{t,i}(\cdot, \boldsymbol{\omega}_{t,i}; Q)$ the value function and continuation function of the SDP, respectively, with hourly decision periods. The SDP formulation becomes

$$V_{t,i}^H(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q) = \max_{(\mathbf{y}_{t,i}, \mathbf{v}_{t,i}, \mathbf{m}_{t,i}) \in \mathcal{Y}_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)} r_{t,i}^H(\mathbf{y}_{t,i}, \boldsymbol{\omega}_{t,i}) + \gamma_2 C_{t,i}(l_{t,i+1}, \boldsymbol{\omega}_{t,i}; Q) \quad (21)$$

$$C_{t,i}(\cdot, \boldsymbol{\omega}_{t,i}; Q) = \mathbb{E}(V_{t,i+1}^H(\cdot, \boldsymbol{\omega}_{t,i+1}; Q) | \boldsymbol{\omega}_{t,i}), \quad (22)$$

with $V_{t,nI}^H(l_{t,nI}, \boldsymbol{\omega}_{t,nI}; Q) = 0$ for each $l_{t,nI}$, and where γ_2 is the weekly discount rate. The immediate reward is defined in (18). The relevant input to the upper level optimization problem in (12) is the discounted accumulated expected revenues from the solution of the resource valuation MDP for a fixed capacity installment Q at time t , i.e. $V_t(\boldsymbol{\omega}_t, q_t(\tau, u)) := V_{t,0}^H(l_{t,0}, \boldsymbol{\omega}_{t,0}; Q)$ with $Q := q_t(\tau, u)$, which fully specifies the optimal stopping and capacity installment MDP in (12). In our numerical section, we consider both the SDP in (21)-(22) which includes short-term operational flexibility, and the SDP in (16)-(17) which excludes short-term operational flexibility in resource valuation.

Computing the continuation function $C_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)$ is intractable. Therefore, we propose in Section 4 approximations which allow us to obtain a feasible policy which we assess against a dual bound (Brown et al. 2010). Before we proceed to the solution approach, we illustrate the added value of including the stage- i intrinsic problem under resource valuation, which better captures the value of storage and capacity to utilize short-term price variations operationally.

3.4 Illustration of resource valuation with and without operations control

We illustrate the short-term aspect by looking at first-stage decisions of problem (16)-(17) and (21)-(22), respectively. We solve each problem near-optimally over a two-year horizon so that we can analyse the first-stage decision for different capacity alternatives and different assumptions regarding within-week prices. This is illustrated in Figure 4. The weekly price profile in year 0 and week 0, $\mathbf{S}_{0,0} = (\exp(\ln S_{t,i} + \alpha_{t,i,j}), j = 1, 2, \dots, 56)$, is illustrated by the solid black line with black points every third hour. The number of periods in the week is the number of hours of the week 168 divided by the number of hours in each block, which is 3 in this example. The dashed line is the weekly price mean. Typically, during the day prices are above the weekly mean and during the

night they are below. During the weekend, prices tend to be below the weekly mean. The bars in the lower panel are near-optimal production decisions in each case. The light shaded grey bars are the first week production decision without short-term flexibility, i.e. $x_{0,0}$, and the dark shaded bars are the first week production decisions with short-term flexibility, i.e. $\mathbf{y}_{0,0} = (y_{0,0,j} \ j = 1, 2, \dots, 56)$. Figure 4a illustrates the case with low production capacity. In this case, the producer generates electricity the whole week if short-term flexibility is ignored, and receives the price 33.4 €/MWh. In terms of the marginal water value (MWV), this means that $MWV < 33.4$ €/MWh. When including short-term flexibility during resource valuation, the producer only produces at high-price periods. In this case, $MWV \approx 35.1$ €/MWh, which is the lowest price the producer generates electricity at in the first week, illustrated by the dotted line in Figure 4a. Figure 4b illustrates the near-optimal first week production decisions when having higher production capacity. Similar to the other case, the producer produces the whole week if short-term flexibility is ignored. If the within-week flexibility is accounted for, the producer requires higher prices for each unit of water produced. In the high capacity case, $MWV \approx 37.1$ €/MWh, which is the dotted line in Figure 4b. This illustrates how capacity and assumptions on within-week prices affect valuation. In the next section, we present our solution approach for the capacity upgrade MDP.

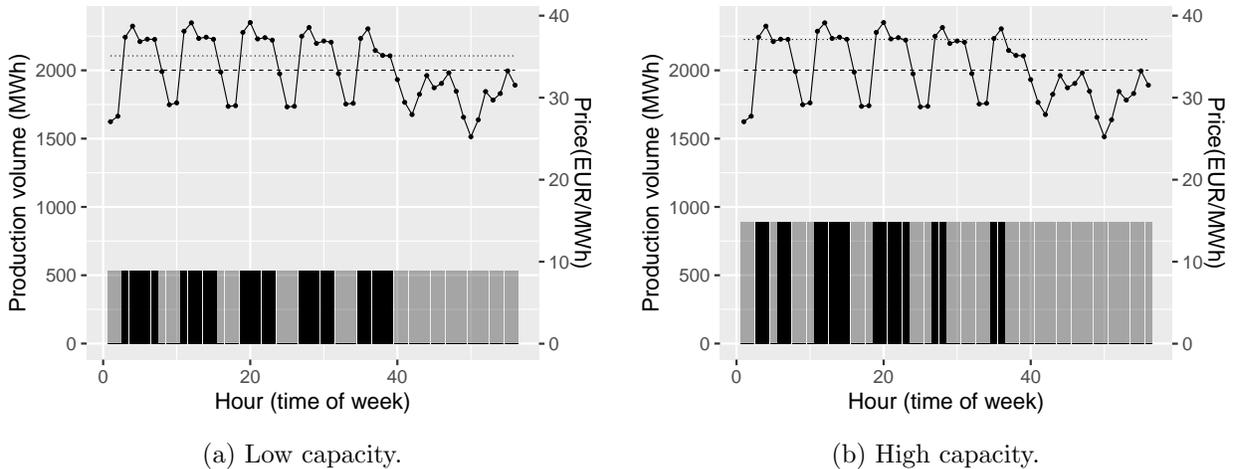


Figure 4: Resource valuation with and without short-term flexibility. The black line with points is the average price per three consecutive hours of the week. The light shaded bars are the near-optimal first week production decisions without short-term flexibility, and the dark shaded bars are the first week production decisions with short-term flexibility. The dashed line is the weekly mean price, and the dotted lines are the marginal water values when incorporating short-term flexibility.

4. Solution approach

This section presents our solution approach for the capacity upgrade MDP in (12). We first impose Assumption 1.

Assumption 1. *Strategic decisions are unaffected by the level of the short-term factors.*

This assumption has been tested empirically in the setting of long-term investments in energy applications using oil futures in Schwartz and Smith (2000). We argue that in our setting, this

assumption is reasonable, since we consider a long planning horizon and short-term effects will quickly diminish. In Appendix D we assess this assumption numerically.

The following proposition characterizes the behaviour of the stage- i value function and continuation function, formulated in (21)-(22), in endogenous states, and is useful in designing a solution approach for the capacity upgrade MDP.

Proposition 4.1. *For a given $(t, i, \boldsymbol{\omega}_{t,i}) \in [0, \infty) \times \mathcal{I} \times \mathbb{R}_+^3$, the functions $V_{t,i}^H(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)$ and $C_{t,i}(l_{t,i}, \boldsymbol{\omega}_{t,i}; Q)$ are concave in $l_{t,i} \in [0, R]$ and $Q \in [q_0, R]$.*

Proof. See Appendix A.1. □

We need to estimate the resource value as a function of capacity. Our approach involves the estimation of the resource value based on a discrete set of capacity alternatives, before employing spline interpolation under the restriction that the overall fit is concave (Pya and Wood 2015), according to Proposition 4.1. Next, we present the estimation of the resource value.

4.1 The reoptimization heuristic

For resource valuation, we apply the reoptimization heuristic (RH), which is popular among practitioners for the management of commodity storage (Lai et al. 2010). An appealing feature of this heuristic is that it straightforwardly can integrate tactical planning and operations control, which is more difficult for other algorithms, such as e.g. stochastic dual dynamic programming, which is common to apply for medium-term hydropower planning (Pereira and Pinto 1991, Gjelsvik et al. 2010). We use RH for obtaining operational policies for different capacity upgrade alternatives. Therefore, we omit in this section the year index t .

Let $F_{i,\nu}$ denote the futures price at time i with time to maturity ν . The vector of hourly futures prices is denoted $\mathbf{F}_{i,\nu} = (F_{i,\nu,j}, j \in \mathcal{J})$. The weekly futures price can be obtained from our price model as the risk-neutral expectation $F_{i,\nu} = \mathbb{E}(S_\nu | \chi_i, \xi_i)$, which is provided in Appendix A.2. Furthermore, let $\bar{Z}_{i,\nu,j} = \mathbb{E}(Z_\nu | \zeta_i)$ and $\bar{\mathbf{Z}}_{i,\nu} = (\bar{Z}_{i,\nu,j}, j \in \mathcal{J})$ denote the weekly inflow mean and the vector of hourly expected inflows at time ν , respectively. The RH policy, $\hat{\pi}$, is obtained by solving intrinsic programs repeatedly. In stage i and state (l_i, ω_i) , the RH model is given by

$$\max_{(\mathbf{y}_\nu, \mathbf{v}_\nu, \mathbf{m}_\nu, \nu \in \mathcal{V}_i)} \sum_{\nu \in \mathcal{V}_i} \sum_{j \in \mathcal{J}} \gamma_2^\nu \gamma_3^j F_{i,\nu,j} y_{\nu,j} \quad (23)$$

$$\text{s.t. } (\mathbf{y}_\nu, \mathbf{v}_\nu, \mathbf{m}_\nu) \in \mathcal{X}_i^{\text{ILP}}(l_i, \omega_i) \quad (24)$$

where $\mathcal{V}_i = \{i, i+1, \dots, nI\}$ is the set of periods from time i to the end of horizon nI , and the decision vector \mathbf{y}_ν contains hourly decisions in week ν . The feasible action set for the stage- i intrinsic program is defined by

$$\mathcal{X}_i^{\text{ILP}}(l_i, \omega_i) = \{$$

$$l_{\nu+1} = l_\nu - \sum_{j \in \mathcal{J}} y_{\nu,j} + \sum_{j \in \mathcal{J}} \bar{Z}_{i,\nu,j} - \sum_{j \in \mathcal{J}} v_{\nu,j} \quad \nu \in \mathcal{V}_i \setminus \{nI\} \quad (25a)$$

$$m_{\nu,j} = l_\nu \quad \nu = i \quad (25b)$$

$$m_{\nu,j+1} = m_{\nu,j} - y_{\nu,j} + \bar{Z}_{i,\nu,j} - v_{\nu,j} \quad \nu \in \mathcal{V}_i, j \in \mathcal{J} \setminus \{J\} \quad (25c)$$

$$m_{\nu,j} - y_{\nu,j} + \bar{Z}_{i,\nu,j} \leq R + v_{\nu,j} \quad \nu \in \mathcal{V}_i, j \in \mathcal{J} \quad (25d)$$

$$y_{\nu,j} \leq m_{\nu,j} + \bar{Z}_{i,\nu,j} \quad \nu \in \mathcal{V}_i, j \in \mathcal{J} \quad (25e)$$

$$y_{\nu,j} \leq \frac{Q}{J} \quad \nu \in \mathcal{V}_i, j \in \mathcal{J} \quad (25f)$$

$$y_{\nu,j}, m_{\nu,j} \geq 0 \quad \nu \in \mathcal{V}_i, j \in \mathcal{J} \quad (25g)$$

}.

Variables and parameters of the feasible action set of the intrinsic program are defined and explained in Section 3.3. The RH policy is defined by repeatedly solving intrinsic programs based on sample paths from the stochastic model. Each time a problem is solved, the first-stage solution of the stage i intrinsic program is implemented. Then the state gets updated, new information gets revealed, and a new intrinsic program gets solved. In our setting, the first-stage decisions are the J first hourly generation decisions. After solving for N sample paths, the lower bound estimate can be found as the sample average of the accumulated discounted revenue by following the RH policy along each sample path,

$$V_0^{\hat{\pi}}(l_0, \boldsymbol{\omega}_0; Q) = \frac{1}{N} \sum_{n=1}^N \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \gamma^i S_{i,j}^n y_{i,j}^n, \quad (26)$$

where $S_{i,j}^n = F_{i,0,j}^n$, is the realized price in stage i and hour j , and $(y_{i,j}^n, j \in \mathcal{J})$ is the implemented first-stage decisions when solving the intrinsic program in stage i and sample n . Proposition 4.2 states a property of the value of the RH policy, which simplifies the computation of the overall policy of the optimal stopping problem in (12).

Proposition 4.2. *The optimal solution of the stage- i intrinsic problem in (23)-(24) is independent of ξ_i .*

Proof. See Appendix A.2. □

Proposition 4.2 entails that the performance of the RH policy can be written as

$$V_0^{\hat{\pi}}(l_0, \boldsymbol{\omega}_0; Q) = e^{\xi_0} G_0(l_0, \chi_0, \zeta_0; Q), \quad (27)$$

Effectively, this means that the tactical and operations control policy, i.e. resource valuation, can be computed for a fixed ξ_0 . The value of the RH policy can then be computed as the product of the long-term level today and the function $G_0(\cdot, \cdot, \cdot; Q)$ which is concave in capacity Q , by Proposition 4.1. Given a representative set of capacity alternatives and the performance of the RH policy for each capacity alternative, we fit a smooth concave function to the data points $(Q_i, G_0(l_0, \chi_0, \zeta_0; Q_k))$, $k \in \{0, 1, \dots, K\}$, where K is a finite number of capacity alternatives. We apply the framework by Pya and Wood (2015) for concave spline interpolation. This function serves then as input to the optimal stopping problem in (12) which we solve analytically. This will be explained next.

4.2 Characterization of the capacity upgrade policy

Under Assumption 1, only the continuous-time movement of the long-term equilibrium price drives the value of the investment. Thus, the value of the option and an investment threshold, ξ^* , above which immediate investment will be optimal can then be derived using Ito's lemma and optimality

conditions (Dixit and Pindyck 1994). As explained in the previous section, we approximate the function $V_t(\boldsymbol{\omega}_t, q_t(\tau, u))$ in (12) by $e^{\xi t} G_t(q_t(\tau, u))$. Under Assumption 1 and Propositions 4.2, the optimal stopping and capacity installment problem in (12) can be written as

$$H(q_0, \xi_0) = \max_{\tau \in [0, \infty), u \in (0, R - q_0]} \mathbb{E} \left[\int_0^\infty \gamma_1^t e^{\xi t} G_t(q_t(\tau, u)) dt - \gamma^\tau K(u) \middle| \xi_0 = \xi \right]. \quad (28)$$

In this representation of the hydropower capacity upgrade problem, only the long-term price enters as an exogenous factor, and the resource value function is separated into the product of the long-term price and a function that is concave in $q_t(\tau, u)$, according to Proposition 4.1. This function represents the annual production output and short-term flexibility which is estimated by the RH policy. Proposition 4.3 characterizes the strategic planning policy and the associated value. The value in the waiting region, $\xi < \xi^*$, consists of two terms. The first term represents the value of potentially increasing capacity if the trigger is met, and the second term is the perpetual revenue of current operations $G(q_0)$ at the current capacity level q_0 . The value in the stopping region, $\xi \geq \xi^*$, is the perpetual operational revenue after installing additional capacity, $G(q_0 + u)$, minus the investment cost $K(u)$. At the time investment occurs, the optimal capacity choice is where the marginal cost of installing capacity equals the marginal capacity value, as stated in (33). Since we do not have an explicit expression for u^* , we solve for D and ξ^* for a fixed u and then evaluate (33). If a convergence criteria is not met, we update our guess for u^* in the direction given by the sign of (33) and a specified magnitude. The magnitude of a move when guessing u^* is halved every time one moves past the solution until the convergence criterion is met. For reasonable starting values, this numerical scheme converges towards the unique optimal solution of (32)-(33).

Proposition 4.3. *The option value in (28) is given by*

$$H(q_0, \xi_0) = \begin{cases} D \exp(\xi)^{\beta_1} + \frac{G(q_0) \exp(\xi)}{\rho} & \text{if } \xi < \xi^*, \\ \frac{G(q_0 + u) \exp(\xi)}{\rho} - K(u) & \text{if } \xi \geq \xi^*. \end{cases} \quad (29)$$

where $\rho = r - (\mu_\xi - \lambda_\xi)$ and

$$D = \frac{1}{\beta_1 \rho \exp(\xi^*)^{(\beta_1 - 1)}} (G(q_0 + u) - G(q_0)), \quad (30)$$

and β_1 is the positive root of

$$\frac{1}{2} \beta(\beta - 1) \sigma_\xi^2 + \beta(\mu_\xi - \lambda_\xi) - r = 0. \quad (31)$$

The investment trigger ξ^* and capacity level u^* are implicitly given by

$$\exp(\xi^*) = \frac{\beta_1}{\beta_1 - 1} \frac{K(u) \rho}{G(q_0 + u) - G(q_0)} \quad (32)$$

$$\frac{\exp(\xi^*)}{\rho} \frac{\partial G(q_0 + u)}{\partial u} \bigg|_{u^*} = \frac{\partial C(u)}{\partial u} \bigg|_{u^*}. \quad (33)$$

Proof. See Appendix A.3 □

4.3 Dual bound estimation

We compute dual bounds for the production policy which determines the resource value. For the dual bound estimation, we relax the non-anticipativity constraints and solve deterministic perfect information problems with dual penalties. Details and theory for dual bounds in stochastic dynamic programs can be found in (Brown et al. 2010). Let $\hat{\omega}^n := \{\hat{\omega}_0^n, \hat{\omega}_1^n, \dots, \hat{\omega}_T^n\}$ denote a vector of realized stochastic variables in each stage $i \in \mathcal{I}$, and where each stage realization consists of J subperiods. We define dual penalties

$$d(\mathbf{y}_\nu, \hat{\omega}^n) = \delta_1 \sum_{\nu \in \mathcal{V}_0} \sum_{j \in \mathcal{J}} y_{\nu,j} \left(\hat{F}_{\nu,j}^n - F_{0,\nu,j} \right) + \delta_2 \sum_{\nu \in \mathcal{V}_0} \sum_{j \in \mathcal{J}} y_{\nu,j} \left(\hat{Z}_{\nu,j}^n - \bar{Z}_{0,\nu,j} \right), \quad (34)$$

where coefficients δ_1 and δ_2 need to be estimated, and the price mean and inflow mean are $F_{0,\nu,j}$ and $\bar{Z}_{0,\nu,j}$, respectively. We restrict the penalty coefficients to be stage-independent. For the estimation, we define a two-dimensional grid and locate approximately which parameter values that lead to the lowest upper bound. Once an area is located, we define a grid of finer granularity and perform a local search up to one significant digit for each coefficient. We estimate the coefficients for a capacity alternative where the gap between the perfect information upper bound and a lower bound obtained by the RH policy is high. We obtain estimates $\delta_1 = -0.8$ and $\delta_2 = -0.00005$. An upper bound can now be attained by solving deterministic problems for N Monte Carlo samples with the estimated dual penalties, and then consider the sample average of revenues obtained from all paths. This bound can then be used to assess the performance of the RH policy, which we demonstrate in our numerical results.

5. Numerical results

In this section, we present the numerical results from a case study of a Norwegian hydropower producer considering upgrading its facilities. We first present the instances before we present results for the capacity upgrade study. We then examine the operational policy and analyse how the operational pattern changes when the facility is upgraded and with assumptions regarding short-term price variations.

5.1 Instances

Plant specific parameters, an estimate of the long-term equilibrium price, and the model parameters of the long-term price dynamics are provided in Table 1. Other parameter estimates that are used in the case study are provided in Appendix B. The reservoir capacity is 335 GWh, and the initial generation capacity and annual average inflow are 8.3% and 24.7% of the reservoir capacity, respectively. This means that about 1/5 of the annual inflow can be stored in the reservoir. Furthermore, the generation station is already sufficiently large to avoid most water spillage in flood periods. We aim to test the impact of changing within-week price variations on investment policies. Therefore, we define in Table 2 three instances. The first instance is the case when using weekly decision periods for resource valuation, $\alpha_{t,i} = \mathbf{0}$. In the second instance, we incorporate the short-term operations control aspect. In numerical experiments we assume that the within week variation is independent of week and year. We estimate the within-week price profile $\alpha_{t,i}$ to historical spot price variations in the period from 2013-2018. We denote this estimate α^{base} . In the

third instance, we scale the 2013-2018 variations by 1.5 to study the effect on investment policies if the within-week price variations increase. We use 3-hourly decision periods, which means that the vector α^{base} has 168/3 elements, similar to the illustrative example in Section 3.4. Estimates for (log) hourly price deviations for all three instances, $\mathbf{0}$, α^{base} , and $1.5\alpha^{\text{base}}$ are plotted in Figure 13 in Appendix B. Figures 4a and 4b illustrate α^{base} on real scale with the weekly mean 33 €/MWh. We run experiments using 2000 samples to compute lower bounds based on the RH policy and upper bounds based on information relaxations and duality theory. The standard errors of the lower and upper bounds are at most 0.53% and 0.24%, respectively.

Table 1: Case study parameters.

		Value	Relative to reservoir capacity
Reservoir capacity	R	335 GWh	1.00
Initial generation capacity	q_0	27.9 GWh/week	0.08
Annual average inflow	$\sum_{i=1}^{52} \bar{\mu}_i$	1354 GWh/year	0.25
Equilibrium price	$\exp(\xi_0)$	30.0 €/MWh	
Long-term drift	$\mu_\xi - \lambda_\xi$	0.012	
Long-term volatility	σ_ξ	0.146	

Table 2: Instances.

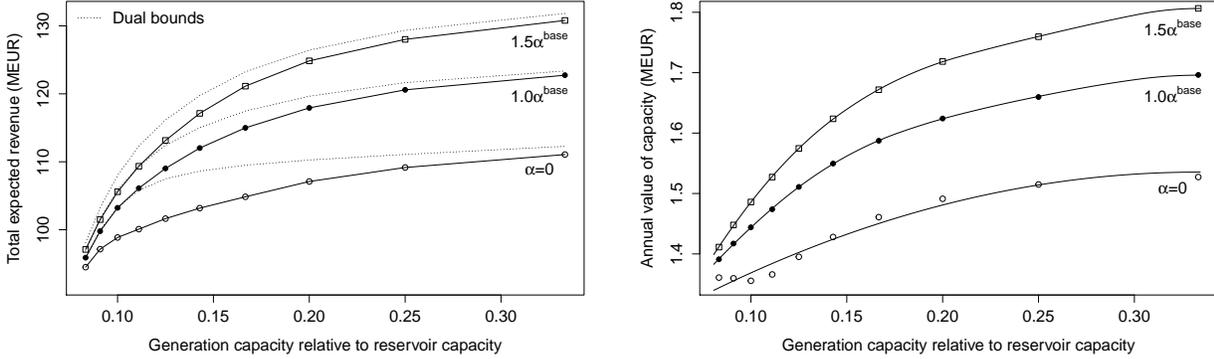
Instance	Description	Spot price data
$\alpha_{t,i} = \mathbf{0}$	Weekly decision periods	Zero variations
$\alpha_{t,i} = 1.0\alpha^{\text{base}}$	3-hourly decision periods	Variations in 2013-2018
$\alpha_{t,i} = 1.5\alpha^{\text{base}}$	3-hourly decision periods	$1.5 \times$ variations 2013-2018

5.2 The capacity upgrade policy

To estimate the cash flow associated with capacity upgrade alternatives, we estimate the resource value $V_0^{\hat{\pi}}(\bar{l}, \bar{\omega}; Q_k)$, where $\hat{\pi}$ denotes the RH policy, under different capacity alternatives Q_k , $k = 1, \dots, K$ ranging from the initial capacity q_0 to one third of the reservoir capacity. The initial reservoir state \bar{l} is the long-run reservoir volume. For details regarding the estimation of \bar{l} , see Appendix B. By Proposition 4.1 we can use an arbitrary long-term level ξ_0 as initial condition for the resource valuation. We use long-term means of short-term factors χ_t and ζ_t , which are zero, as the initial conditions for the short-term factors. By doing so, we get an estimate of the annual cashflow for each capacity alternative for a representative year and long-term equilibrium price level.

Figure 5a shows the performance of the RH policies. The dotted lines are dual upper bounds for each instance $\alpha = \mathbf{0}$, α^{base} , and $1.5\alpha^{\text{base}}$. The leftmost point, i.e. generation capacity $q_0 = 0.08$ relative to the maximum reservoir capacity, is the RH policy performance without upgrades. The solid curves show the performance for a set of generation capacity upgrade alternatives between 0.08 and 0.33 of the reservoir capacity. We observe that the total expected accumulated revenue improves when including operations control, and is higher when within-week price variations are higher. Figure 5b shows the performance of the RH-policy for a representative year relative to the long-term price $\exp(\xi_0)$. That is, the figure shows $G(q_0+u)$ as a function of q_0+u for $u \in [0, 0.33-q_0]$.

The lines are spline interpolations under the restriction that the overall fit is concave, according to Proposition 4.1. These are estimated using the approach by Pya and Wood (2015). For consistent resource valuation, we make sure that the initial and the average end of year 1 reservoir volume is approximately equal. For details, see Appendix B.



(a) Performance of the RH-policy for two-year production scheduling. (b) Estimated concave functions of annual revenue. The x-axis is $q_0 + u$ and the y-axis is $G(q_0 + u)$.

Figure 5: RH policy performances as a function of generation capacity.

Table 3 highlights the effect of within-week price assumptions on capacity installments. The net present value (NPV) of continuing operations at the current capacity and current long-term price $\exp(\xi_0) = 30$ €/MWh is provided in the first row. By comparing this with the second row where the investment occurs immediately, it is clear that the project should be undertaken in all instances according to the NPV criterion. The third row reports the real options value, which takes into account the opportunity cost of realizing the project. Capacity installments are reported in the next two rows. The fourth row reports the capacity installments if investing now, and the fifth row reports the capacity installment at the price trigger. The price trigger is reported in the next row. Only the case with high within-week price variations invests immediately, since the trigger is $\exp(\xi^*) = 22.0 < 30.0 = \exp(\xi_0)$. We observe that under $1.5\alpha^{\text{base}}$ one will invest immediately in 87.6% additional capacity, while under α^{base} one will invest in a lower capacity, 83.7%, with 85.0% probability in the next 10 years, as seen in the last row. Note that capacity installments are reported as a percentage of the current capacity q_0 . Compared to the case $\alpha = 0$ one will invest earlier and in significantly more capacity at a given price. The optimal capacity upgrade at the price trigger is higher in $\alpha = 0$ compared to the other cases since the price trigger is much higher, which means that the investment will be undertaken only with 1.3% probability over the next 10 years.

Figure 6 shows sensitivity of the investment trigger in σ_ξ . The areas illustrate at which price it is optimal to undertake the investment project. If the current price is below the light grey shaded area, it is optimal to wait in all instances. If it is in the light gray area, it is optimal to invest only in the $1.5\alpha^{\text{base}}$ instance. In the grey area one will invest in both instances α^{base} and $1.5\alpha^{\text{base}}$, and in the dark grey area one will invest in all instances. Investment triggers behave as expected, in the sense that increased uncertainty increases the value of waiting and hence the price trigger

increases.

Table 3: The investment policy and value under different assumptions regarding within week price variations: 0, 1 and 1.5 times the variation the most recent 5 years. Values are reported in million €. Additional capacity installments are reported as a percentage of current capacity q_0 .

		$\alpha = 0$	$1.0\alpha^{\text{base}}$	$1.5\alpha^{\text{base}}$
NPV (no investment)		5 039.7	5 219.0	5 294.8
NPV (investment)		5 231.5	5 711.9	6 009.3
Real options value	$H(q_0, \xi_0)$	5 256.3	5 712.1	6 009.3
Capacity upgrade (invest now)	$u^*(\xi_0)$	71.3%	82.3%	87.6%
Capacity upgrade (at trigger)	$u^*(\xi^*)$	101%	83.7%	80.4%
Price trigger	$\exp(\xi^*)$	93.3 €/MWh	32.0 €/MWh	22.0 €/MWh
Investment probability (10 years)		1.3%	85.0%	100%

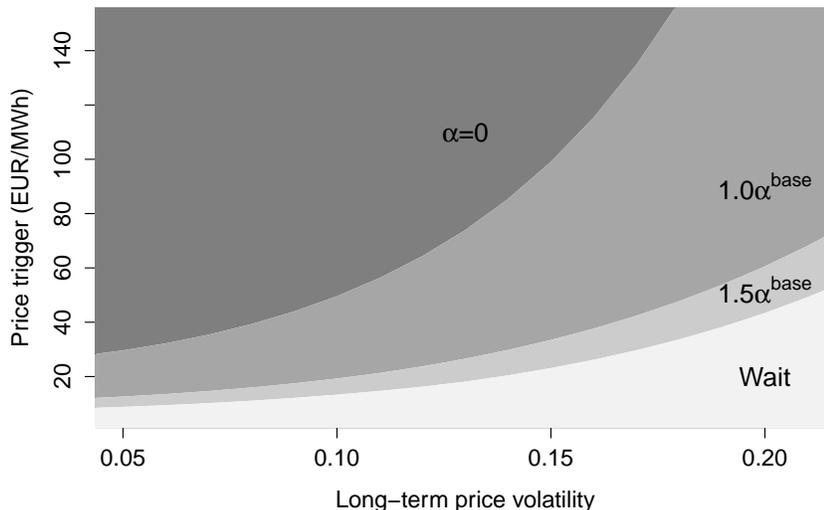


Figure 6: Price trigger sensitivity. The shaded shows at which prices it is optimal to invest in each of the instances.

Figure 7 shows the additional value generated by investing in u units of capacity at the price trigger of the α^{base} instance, $\exp(\xi^*) = 32$ €/MWh, together with the cost of investing in capacity u . The optimal capacity installment in each instance is where the marginal value of installing capacity equals the marginal cost.

5.3 The operational policy

In this section, we take a closer look at operational policies, which determines the cashflow associated with capacity upgrade alternatives. Figure 8 illustrates the reservoir trajectories when generation is constrained by the initial maximum generation capacity q_0 , and with different assumptions regarding decision periods. The figures show that the behaviour in all instances is similar. This is because the operational flexibility is limited. However, we observe that the per-

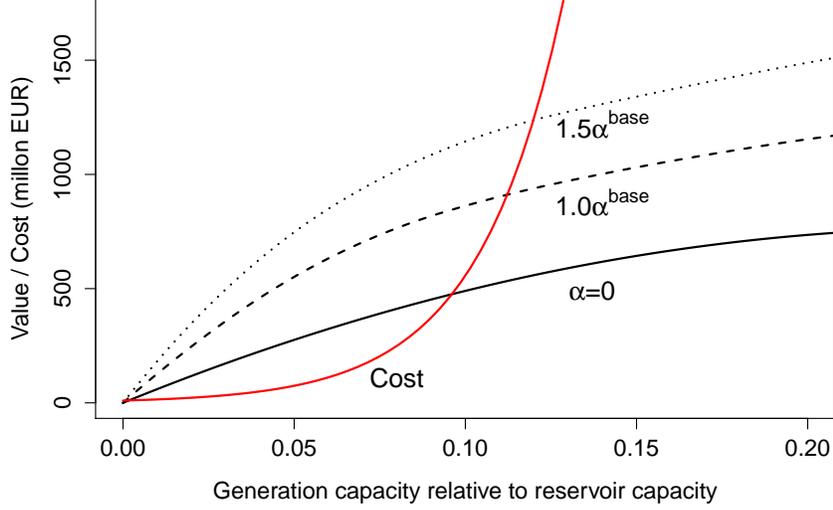


Figure 7: Value of additional capacity $\frac{\exp(\xi)}{\rho} (G(q_0 + u) - G(q_0))$ at price $\exp(\xi) = 32$, which is the price trigger for the instance with α^{base} , as a function of capacity installment u . The red line is the investment cost $K(u) = A \exp(Bu)$.

centile bands are slightly broader in the instances α^{base} and $1.5\alpha^{\text{base}}$ where the operations control aspect is considered. Figure 9 shows the reservoir trajectories after a 70% capacity increase. In this case, the differences are more prominent. For instance $\alpha = \mathbf{0}$ in Figure 9a we observe that the producer prefers to have a full reservoir during the fall when the price expectation tends to be higher than other parts of the year. For α^{base} and $1.5\alpha^{\text{base}}$ in Figure 9b and 9c, respectively, the reservoir trajectories show less variation and broader percentile bands as opposed to the case when the operations control aspect is neglected. The effect can be explained as follows: In the instance $\alpha = \mathbf{0}$ the producer values its water more in the future during weeks 35-45, since prices are expected to increase due to seasonality effects and the weekly operational flexibility is sufficiently high after the capacity increase. Thus, the producer chooses to produce nothing, or only the incoming inflow during these weeks. In the instances α^{base} and $1.5\alpha^{\text{base}}$, the producer has within-week flexibility to produce in parts of the week when prices are high, e.g. during daytime. Therefore, the producer uses part of the water stored in the reservoir in most weeks throughout weeks 35-45, which leads to a lower expected reservoir trajectory and a broader percentile band.

Figure 10 shows the average generation and average spillage for each of the three instances over the two-year horizon used for the resource valuation. The graphs in Figure 10a illustrate that the producer can not increase the production output much by investing in additional generation capacity, only from about 8.2 to 9.0 over two years. The numbers are presented relative to the reservoir capacity. This means that the additional value from capacity installments stems mostly from the flexibility to exploit price variations. Figure 10b shows that the average spillage decreases with capacity. The highest spillage happens in the instance $\alpha = \mathbf{0}$, while the instances α^{base} and $1.5\alpha^{\text{base}}$ have less spillage since the reservoir is managed differently, as explained above.

Figure 11 shows optimality gaps for each instance. The solid lines are the percentage difference

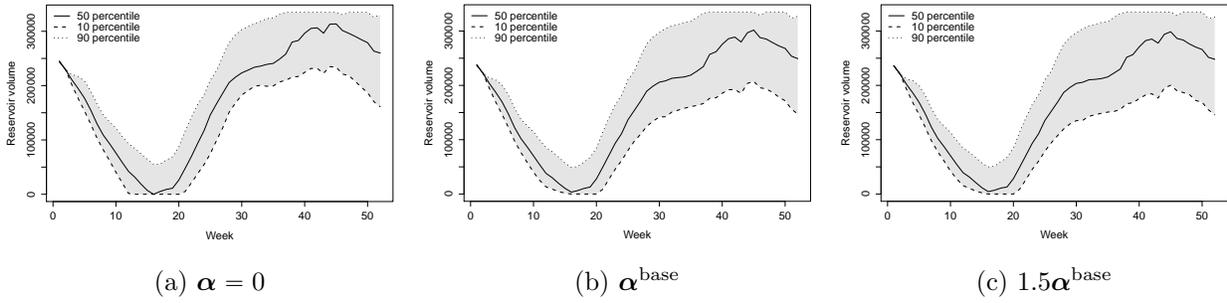


Figure 8: Reservoir trajectories before investment with maximum capacity q_0 .

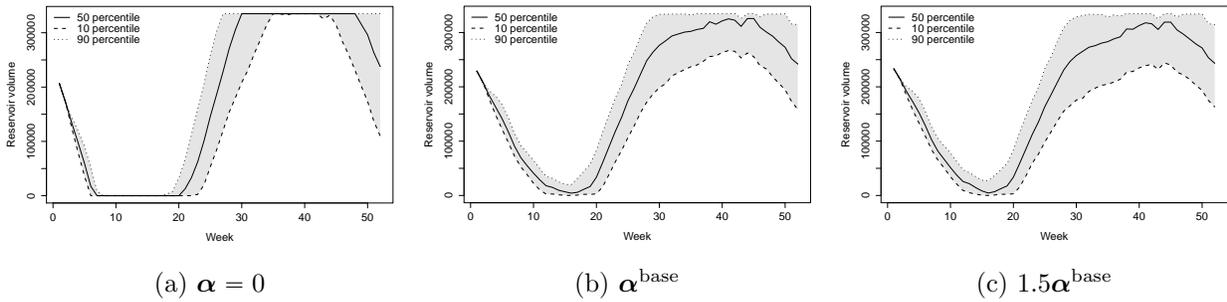


Figure 9: Reservoir trajectories after investment in maximum capacity $u = 0.70q_0$.

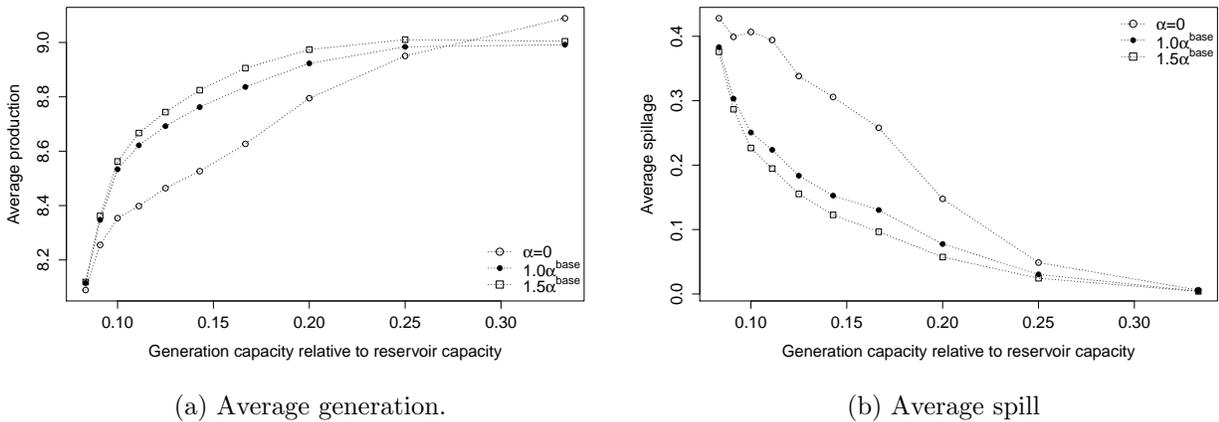


Figure 10: Total average generation and spill as a function of capacity alternatives. The curves are plotted relative to the maximum reservoir capacity.

between the upper bound with estimated dual penalties and the lower bound, i.e. the value of the RH policy in each instance. The dotted lines are the gap compared to perfect information upper bounds without dual penalties. We observe that the optimality gap based on the dual bound is lower for the instances α^{base} and $1.5\alpha^{\text{base}}$. This can partly be explained by hourly variations in prices which makes the RH-policy less aggressive towards the upper reservoir limit, and thus manages spillage risk better. The instance $\alpha = 0$ obtains an optimality gap of at most 5.5%

from the optimal policy value. The instances α^{base} and $1.5\alpha^{\text{base}}$ obtain at most 3.0% and 2.6%, respectively.

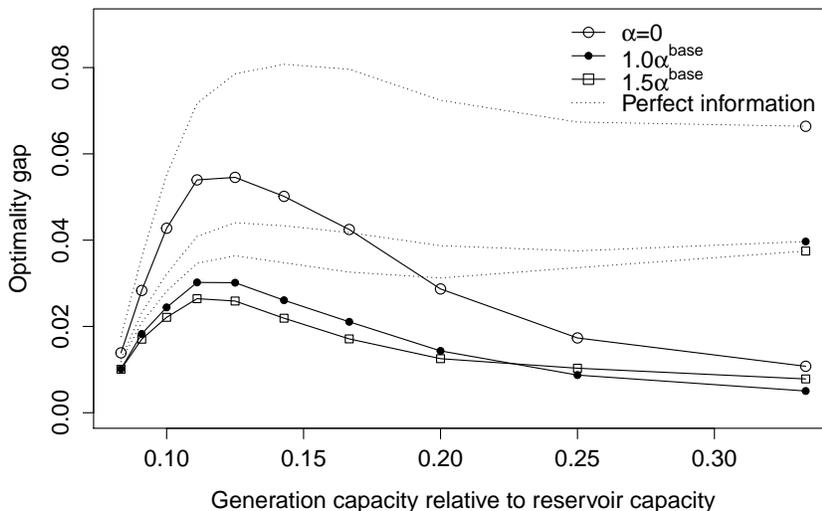


Figure 11: Optimality gap for each instance using dual penalties $\delta_1 = -0.8$ and $\delta_2 = -0.00005$. The dotted lines are perfect information bounds with zero dual penalty.

6. Conclusions

Accurate mathematical models for the calculation of cashflows associated with investments in hydropower are important. Hydropower plants with storage reservoirs are flexible and can quickly ramp up and down production. Consequently, hydropower producers can exploit hourly variations in prices, which can support the integration of intermittent energy sources in the ongoing transition to a more sustainable power system. Moreover, since weather and prices are uncertain, hydropower producers need to plan production over several years to estimate their marginal water values. This value is the expected price at which power generated from a unit of water can be sold for in the spot market if the timing of this sale is optimized. In addition to weather and market dynamics, the resource valuation depends on reservoir and generation capacity constraints. Therefore, the valuation of generation capacity upgrades in hydropower plants with reservoirs is a complex optimization problem. To handle the complexity of the problem, we revisit hierarchical planning. We formulate the hierarchical planning problem as an MDP. We show how a combination of price modeling, informed by empirical analysis, the MDP structure, and we use reinforcement learning for operations planning can lead to insightful semi-analytical policies. Our analysis demonstrates that assumptions on short-term price variations have a significant impact on the valuation of investment projects in hydropower. Furthermore, our case study shows that if the short-term flexibility of the hydropower producer is accounted for, the net present value of plant operations increases, while a higher capacity is expected to be installed sooner.

Acknowledgments

The first and third authors gratefully acknowledge support from the research center HydroCen, RCN No. 257588. The second author appreciates research support from the College of Business at the University of Illinois at Chicago. We are thankful to Montel for supplying price data, and to Eviny for supplying data regarding their hydropower plants.

References

- Andersson, A. M., Elverhøi, M., Fleten, S.-E., Fuss, S., Szolgayová, J., and Troland, O. C. (2014). Upgrading hydropower plants with storage: timing and capacity choice. *Energy Systems*, 5(2):233–252.
- Anthony, R. N. (1965). *Planning and control systems: a framework for analysis*. MA: Harvard Business School, Boston.
- Arvanitidis, N. V. and Rosing, J. (1970). Composite representation of a multireservoir hydroelectric power system. *IEEE Transactions on Power Apparatus and Systems*, (2):319–326.
- Belsnes, M. M., Wolfgang, O., Follestad, T., and Aasgård, E. K. (2016). Applying successive linear programming for stochastic short-term hydropower optimization. *Electric Power Systems Research*, 130:167–180.
- Benth, F. E., Koekkebakker, S., and Ollmar, F. (2007). Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *The Journal of Derivatives*, 15(1):52–66.
- Bitran, G. R. and Tirupati, D. (1993). Hierarchical production planning. *Handbooks in Operations Research and Management Science*, 4:523–568.
- Bøckman, T., Fleten, S.-E., Juliussen, E., Langhammer, H. J., and Revdal, I. (2008). Investment timing and optimal capacity choice for small hydropower projects. *European Journal of Operational Research*, 190(1):255–267.
- Boomsma, T. K., Meade, N., and Fleten, S.-E. (2012). Renewable energy investments under different support schemes: A real options approach. *European Journal of Operational Research*, 220(1):225–237.
- Borghetti, A., D’Ambrosio, C., Lodi, A., and Martello, S. (2008). An milp approach for short-term hydro scheduling and unit commitment with head-dependent reservoir. *IEEE Transactions on Power Systems*, 23(3):1115–1124.
- Brown, D. B., Smith, J. E., and Sun, P. (2010). Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4-part-1):785–801.
- Conejo, A. J., Arroyo, J. M., Contreras, J., and Villamor, F. A. (2002). Self-scheduling of a hydro producer in a pool-based electricity market. *IEEE Transactions on Power Systems*, 17(4):1265–1272.
- Dangl, T. (1999). Investment and capacity choice under uncertain demand. *European Journal of Operational Research*, 117(3):415–428.
- Dempster, M. A., Fisher, M., Jansen, L., Lageweg, B., Lenstra, J. K., and Rinnooy Kan, A. (1981). Analytical evaluation of hierarchical planning systems. *Operations Research*, 29(4):707–716.
- Dimoski, J., Fleten, S.-E., Löhndorf, N., and Nersten, S. (2019). Dynamic hedging for the real option management of hydropower production with exchange rate risks. Working paper.
- Dixit, R. K. and Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press.
- Duffie, D. (2010). *Dynamic asset pricing theory*. Princeton University Press.
- EIA (2017). Hydroelectric generators are among the United States’ oldest power plants - Today in Energy - U.S. Energy Information Administration (EIA). Retrieved from <https://www.eia.gov/todayinenergy/detail.php?id=30312>. Accessed December 05, 2019.

- Flatabø, N., Haugstad, A., Mo, B., and Fosso, O. B. (1998). Short-term and medium-term generation scheduling in the norwegian hydro system under a competitive power market structure. In *EPSOM'98 (International Conference on Electrical Power System Operation and Management)*, Switzerland.
- Fleten, S.-E., Maribu, K. M., and Wangensteen, I. (2007). Optimal investment strategies in decentralized renewable power generation under uncertainty. *Energy*, 32(5):803–815.
- Gjelsvik, A., Mo, B., and Haugstad, A. (2010). Long-and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming. *Handbook of Power Systems I*, pages 33–55.
- Goodwin, D. (2020). Schwartz-Smith 2-factor model - Parameter estimation. Retrieved from <https://www.mathworks.com/matlabcentral/fileexchange/43352-schwartz-smith-2-factor-model-parameter-estimation>. Accessed June 23, 2021.
- Hagspiel, V., Huisman, K. J., and Kort, P. M. (2016). Production flexibility and capacity investment under demand uncertainty. *International Journal of Production Economics*, pages 95–108.
- Huisman, K. J. and Kort, P. M. (2015). Strategic capacity investment under uncertainty. *The RAND Journal of Economics*, 46(2):376–408.
- IRENA (2015). Hydropower Technology Brief.
- Kleiven, A., Risanger, S., and Fleten, S.-E. (2021). Co-movements between forward prices and resource availability in hydro-dominated electricity markets. Working paper.
- Lai, G., Margot, F., and Secomandi, N. (2010). An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research*, 58(3):564–582.
- Lenstra, J. K., Kan, A. R., and Stougie, L. (1984). A framework for the probabilistic analysis of hierarchical planning systems. *Annals of Operations Research*, 1(1):23–42.
- Liu, L., Zhang, M., and Zhao, Z. (2019). The application of real option to renewable energy investment: A review. *Energy Procedia*, 158:3494–3499.
- Löhndorf, N. and Wozabal, D. (2021). Gas storage valuation in incomplete markets. *European Journal of Operational Research*, 288(1):318–330.
- Löhndorf, N., Wozabal, D., and Minner, S. (2013). Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research*, 61(4):810–823.
- Lucia, J. J. and Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1):5–50.
- Montel (2021). Montel.
- Nadarajah, S. and Secomandi, N. (2018). Merchant energy trading in a network. *Operations Research*, 66(5):1304–1320.
- Nadarajah, S. and Secomandi, N. (2021). Real options in energy: A guided analysis of the operations literature. *Available at SSRN 3877512*.
- NordPool (2021). Nordpool. Retrieved from <https://www.nordpoolgroup.com/>. Accessed August 24, 2021.
- Pereira, M. V. and Pinto, L. M. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1-3):359–375.
- Pérez-Díaz, J. I., Guisández, I., Chazarra, M., and Helseth, A. (2020). Medium-term scheduling of a hydropower plant participating as a price-maker in the automatic frequency restoration reserve market. *Electric Power Systems Research*, 185:106399.
- Philpott, A. B. and Guan, Z. (2008). On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters*, 36(4):450–455.
- Pyra, N. and Wood, S. N. (2015). Shape constrained additive models. *Statistics and Computing*, 25(3):543–559.

- Schwartz, E. and Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911.
- Secomandi, N. (2015). Merchant commodity storage practice revisited. *Operations Research*, 63(5):1131–1143.
- Shapiro, A., Tekaya, W., da Costa, J. P., and Soares, M. P. (2013). Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research*, 224(2):375–391.
- Wallace, S. W. and Fleten, S.-E. (2003). Stochastic programming models in energy. *Handbooks in Operations Research and Management Science*, 10:637–677.
- Wu, O. Q., Wang, D. D., and Qin, Z. (2012). Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management*, 14(3):455–471.

A. Proof of Propositions

A.1 Proof of Proposition 4.1

Proof. We omit the index t in the proof. No production results in zero value, hence $V_i^H(l_i, \boldsymbol{\omega}_i; Q) \geq 0$, which implies $C_i(l_{i+1}, \boldsymbol{\omega}_i; Q) \geq 0$. An upper bound is given by the maximum production in every period:

$$\begin{aligned} C_i(l_{i+1}, \boldsymbol{\omega}_i; Q) &= \gamma_2 \mathbb{E} [V_{i+1}^H(l_{i+1}, \boldsymbol{\omega}_{i+1}; Q) | \boldsymbol{\omega}_i] \\ &\leq \frac{Q}{J} \left(\sum_{\nu=\mathcal{V}_i} \sum_{j=\mathcal{J}} \gamma_2^\nu \gamma_3^j F_{i,\nu,j} \right) \\ &< \infty, \end{aligned}$$

where the set \mathcal{V}_i is the set of time periods from stage i to I , and where the last inequality follows by the martingale property of futures prices. Thus, the value function and continuation function are finite. At stage $I-1$ the continuation function is zero and thus concave. The value function at stage $I-1$ is given by

$$V_{I-1}^H(l_{I-1}, \boldsymbol{\omega}_{I-1}; Q) = \max_{(\mathbf{y}_{I-1}, \mathbf{v}_{I-1}, \mathbf{m}_{I-1}) \in \mathcal{Y}_{I-1}(l_{I-1}, \boldsymbol{\omega}_{I-1}; Q)} r_{I-1}^H(\mathbf{y}_{I-1}, \boldsymbol{\omega}_{I-1}) \quad (35)$$

This is a linear program where l_{I-1} and Q are upper bounds on the convex feasible action set $\mathcal{Y}_{I-1}(l_{I-1}, \boldsymbol{\omega}_{I-1}; Q)$. Thus, concavity of $V_{I-1}^H(\cdot, \boldsymbol{\omega}_{I-1}; \cdot)$ follows from standard linear programming results. By finiteness of the continuation function and the induction hypothesis, $C_i(\cdot, \boldsymbol{\omega}_i; \cdot)$ is concave in $l_{i+1} \in [0, R]$ and $Q \in (0, R - q_0]$. It follows that for a convex set $\mathcal{Y}_i(l_i, \boldsymbol{\omega}_i; Q)$ bounded by l_i and Q , and a linear revenue function $r_i^H(\mathbf{y}_i, \boldsymbol{\omega}_i)$, the value function $V_i^H(\cdot, \boldsymbol{\omega}_i; \cdot)$ is concave. \square

A.2 Proof of Proposition 4.2

Proof. The objective in (23) can be written as

$$\sum_{\nu \in \mathcal{V}_i} \sum_{j \in \mathcal{J}} \gamma_2^\nu \gamma_3^j F_{i,\nu,j} y_{\nu,j} = \sum_{\nu \in \mathcal{V}_i} \sum_{j \in \mathcal{J}} \gamma_2^\nu \gamma_3^j \mathbb{E}(S_\nu | \xi_i, \chi_i) y_\nu \quad (36)$$

$$= \sum_{\nu \in \mathcal{V}_i} \sum_{j \in \mathcal{J}} \gamma_2^\nu \gamma_3^j e^{\xi_i + e^{-\kappa\nu} \chi_i + L(i,\nu)} y_\nu \quad (37)$$

$$= e^{\xi_i} \left(\sum_{\nu \in \mathcal{V}_i} \sum_{j \in \mathcal{J}} \gamma_2^\nu \gamma_3^j e^{-\kappa\nu \chi_i + L(i,\nu)} x_\nu \right), \quad (38)$$

where $L(i, \nu)$ are time-dependent parameters in the risk-neutral price expectation expression. The expression for futures prices traded at time i with maturity at time ν is given by the expected price in (1) under the risk-neutral measure,

$$F_{i,\nu} = \mathbb{E}_Q(S_\nu | \chi_i, \xi_i) \quad (39)$$

$$= \exp \left(e^{-\kappa\chi\nu} \chi_i + \xi_i + (\mu_\xi - \lambda_\xi)\nu - (1 - e^{-\kappa\chi\nu}) \frac{\lambda_\chi}{\kappa\chi} \right) \quad (40)$$

$$+ \frac{\sigma_\chi^2}{2\kappa_\chi} (1 - e^{-\kappa_\chi \nu}) + (1 - e^{-\kappa_\chi \nu}) \frac{\rho \chi \xi \sigma_\chi \sigma_\xi}{\kappa_\chi} \Big). \quad (41)$$

Thus, since $e^{\xi_i} > 0$ in (38), the optimal solution of the stage- i intrinsic program is independent of ξ_i . \square

A.3 Proof of Proposition 4.3

Proof. In this proof, we define $p = \exp(\xi)$ for notational convenience. We derive the option value and investment trigger p^* for a fixed u . In the stopping region, $p > p^*$, the value is given by

$$H(q_0, p; u) = \mathbb{E} \left[\int_0^\infty \gamma^t p_t G_t(q_0 + u) dt - K(u) \Big| \xi_0 = \xi \right] = \frac{G(q_0 + u)p}{\rho} - K(u). \quad (42)$$

In the continuation region, $p < p^*$, Bellman's equation must hold. We obtain the differential equation

$$rH - \frac{1}{2}\sigma_\xi^2 p^2 H''(p) - (\mu_\xi - \lambda_\xi)pH'(p) - pG(q_0) = 0. \quad (43)$$

The general solution of this equation is of the form

$$H(p) = Dp^{\beta_1} + Ep^{\beta_2} + \Psi(p), \quad (44)$$

where β_1 and β_2 are the positive and negative roots of the fundamental equation

$$\frac{1}{2}\beta(\beta - 1)\sigma_\xi^2 + \beta(\mu_\xi - \lambda_\xi) - r = 0. \quad (45)$$

The negative root must be disregarded to prevent the value to become infinitely large when the price approaches zero. Thus, we need $E = 0$. The particular solution takes the form

$$\Psi(p) = \frac{G(q_0)p}{\rho}, \quad (46)$$

where $\rho = r - \mu_\xi - \frac{1}{2}\sigma_\xi^2$, which can be verified by inserting this into (43). Thus, the value in the continuation region is of the form

$$H(p) = Dp^\beta + \frac{G(q_0)p}{\rho}. \quad (47)$$

In addition, H must satisfy

$$H(0) = 0 \quad (48)$$

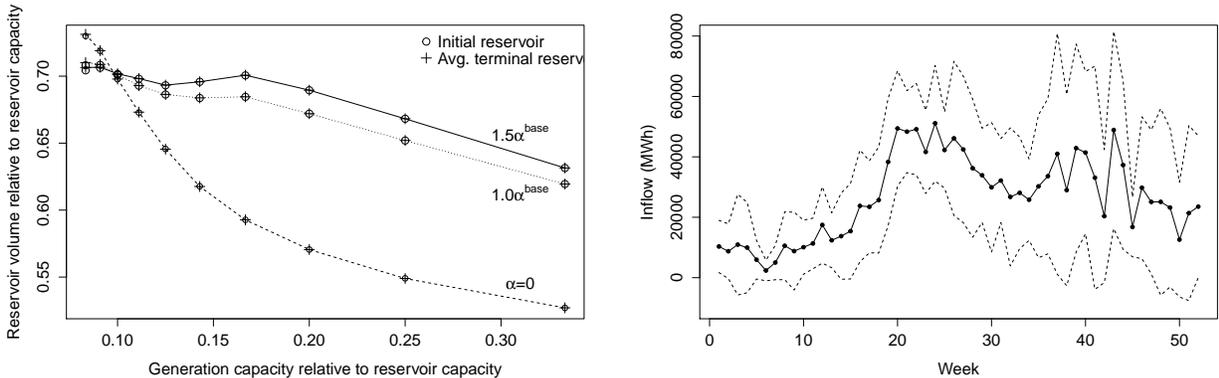
$$H(p^*) = \frac{G(q_0 + u)p^*}{\rho} - K(u) \quad (49)$$

$$H'(p^*) = \frac{G(q_0 + u) - G(q_0)}{\rho} \quad (50)$$

Using (47) and solving for D and p^* , we get the expressions in Proposition 4.3. The optimal capacity choice can now be found by maximizing (42) with respect to the capacity installment u . Thus, the first-order condition gives the optimal capacity choice. \square

B. Parameter estimates

This section provides a description of the parameters and parameter estimates used in Section 5. For calibration of the inflow model, we use maximum likelihood estimation. Inflow data are provided by the producer at the location we consider in the case study. For model calibration, we collect spot data and futures contracts from Montel (2021) using futures contracts traded in the period 2013-2018 with maturities ranging from 1 to 60 months from a synthetic forward curve computed based on average-based forward contracts Benth et al. (2007). Furthermore, we get hourly spot price data from NordPool (2021). The calibration was done using Kalman filtering and maximum likelihood estimation (Goodwin 2020).



(a) Initial reservoir as a function of generation capacity. (b) Inflow mean and standard deviation for each week of a year.

Figure 12: Reservoir and inflow.

Figure 12a compares the starting reservoir volume with the expected end reservoir volume when using the expected end reservoir volume (from the solution of an initial run) as the starting reservoir. We observe that the initial start and expected end reservoir are approximately the same for each run and all capacity alternatives. Therefore, we conclude that one year of operations is sufficient in capturing the steady-state operations. Although different instances have different starting reservoirs, all capacity valuations use approximately the same amount of water in expectation since the expected end reservoir approximately maps the initial reservoir. This gives a fair value of steady-state operations for all capacity alternatives.

Figure 12b shows the mean and standard deviation of inflow for each week of the year. We observe that there are strong seasonal effects. There is an inflow peak during the spring, while the inflow variance is highest during the fall. Figure 13 shows the average deviation from the weekly mean for each hour of the week based on spot price data for the period 2012-2018.

Table 4: Parameters used in the case study.

State variables		
l_t [MWh]	Reservoir volume	
χ_t [€/MWh] (log)	Price deviation from long-term price	
ξ_t [€/MWh] (log)	Long-term price	
ζ_t [MWh]	Inflow deviation from the mean	
Values derived from states		
Z_t [MWh]	Local inflow	
S_t [€/MWh]	Spot price	
Decision variables		
x_t and \mathbf{y}_t [MWh/period]	Production	
z_t and \mathbf{v}_t [MWh/period]	Spillage	
\mathbf{m}_t [MWh]	Reservoir constraints	
q_t [MWh/week]	Maximum production capacity	
Parameters		Estimate
r	risk-free rate	0.02
R [MWh]	Maximum reservoir volume	334 989
A	Cost parameter	10 000 000
B	Cost parameter	0.00012
ϕ_1 [€/MWh] (log)	Price seasonality sine parameter	-0.025
ϕ_2 [€/MWh] (log)	Price seasonality parameter, shift	0.163
κ_χ [1]	Short-term price mean reversion	1.217
λ_χ [€/MWh] (log)	Short-term price risk premium	0.042
$\mu_\xi - \lambda_\xi$ [€/MWh] (log)	Long-term risk-adjusted drift	0.012
$\rho_{\chi\xi}$ [1]	Price factor correlation	0.034
κ_ζ [1]	Inflow mean reversion	52.0
σ_χ [MWh] (log)	Volatility short-term price	0.467
σ_ξ [MWh] (log)	Volatility long-term price	0.146
σ_ζ [MWh]	Standard deviation inflow	6.32
$\bar{\mu}_t, t = 1, \dots, 52$ [MWh]	Weekly historical local inflow mean	Figure 12b
$\bar{\sigma}_t, t = 1, \dots, 52$ [MWh]	Weekly historical local inflow standard deviation	Figure 12b
$\alpha_{t,i,j}, j = 1, \dots, 56$ [€/MWh]	Average price pr 3 consecutive hours of the week	Figure 13
Initial states		
l_0 [MWh]	Initial reservoir	Figure 12a
q_0 [MWh/week]	Initial capacity	27 916
χ_0 [€/MWh]	Initial price deviation	0
ζ_0 [€/MWh]	Initial inflow deviation	0
$\exp(\xi_0)$ [€/MWh]	Initial equilibrium price	30.0

C. Numerical assessment of price dynamics

To assess the assumptions on price dynamics by applying the model in [Schwartz and Smith \(2000\)](#) to electricity price dynamics, we compare the model fit of our geometric Gaussian price model with an alternative additive Gaussian price model, which has been proposed as alternative price dynamics ([Lucia and Schwartz 2002](#), [Kleiven et al. 2021](#)). The additive model can be written as

$$S_t^A = \phi_1^A \cos\left((t + \phi_2^A) \frac{2\pi}{k}\right) + \chi_t^A + \xi_t^A, \quad (51)$$

$$d\chi_t^A = -\kappa_\chi^A \chi_t^A dt + \sigma_\chi^A dz_\chi^A \quad (52)$$

$$d\xi_t^A = \mu^A dt + \sigma_\xi^A dz_\xi^A, \quad (53)$$

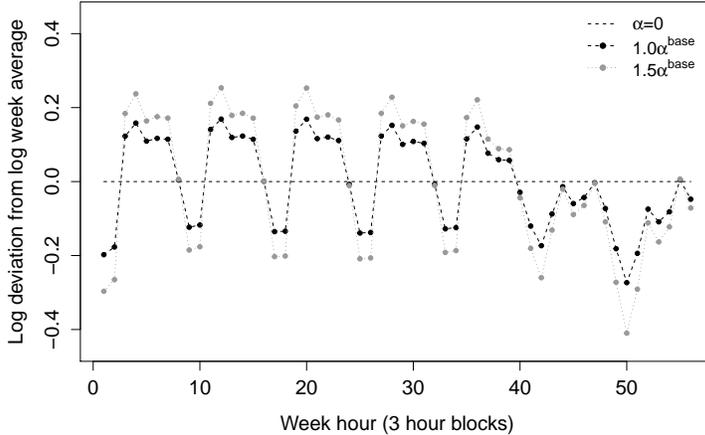


Figure 13: Hourly log price deviations from the log weekly mean, averaged over 3 consecutive hours. For weekly decision periods $\alpha = \mathbf{0}$. The estimated average deviations in the period 2013-2018 is α^{base} . The grey stippled line illustrates $1.5\alpha^{base}$.

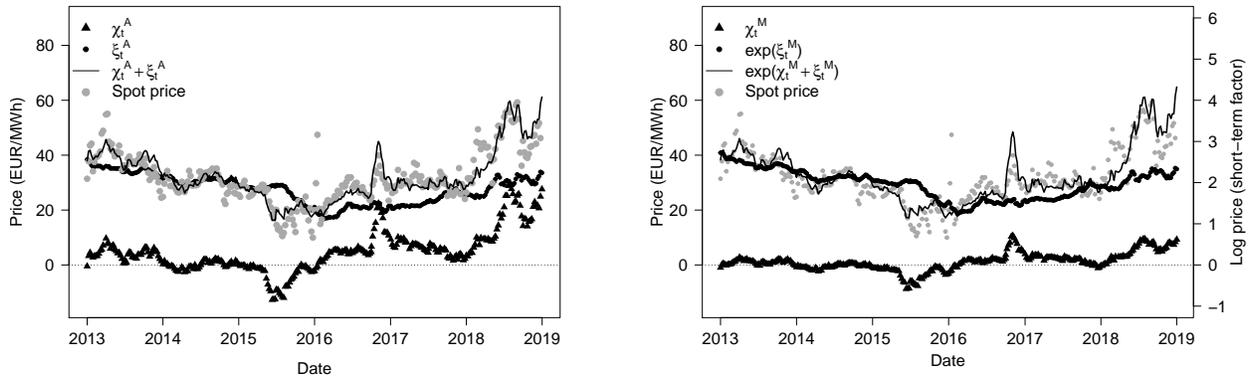
where A abbreviates *additive*. In this section we assess the model fit of the additive model and compare it to the logarithmic model in (2)-(3). We follow Schwartz and Smith (2000) and Kleiven et al. (2021) for the derivation of risk-neutral multiplicative and additive processes, respectively.

In our setting, we search for a model that can capture futures curve dynamics and deterministic time of week spot price behaviour. Thus, we need to assess the fit to futures contracts and hourly spot prices. Similarly, for the price model we use, we define the additive model $S_t^A + \alpha_t^A$ for hourly prices, where S_t^A is the week mean price and α_t^A accounts for time-of-week effects.

Latent state estimates for the additive and multiplicative model in (2)-(3) are plotted in Figure 14a and Figure 14b, respectively, together with the weekly average spot price. In Figure 14b short-term deviations are plotted on log scale, while the long-term equilibrium level is exponentiated and thus plotted on real scale. We observe that the lines $\chi_t^A + \xi_t^A$ in Figure 14a and $\exp(\chi_t^M + \xi_t^M)$ are close to each other, both providing good approximations of the spot price F_t . Figure 15a shows hourly spot prices together with the weekly average price and the estimated latent states from both models. Figure 15b shows the average spot price deviation from the weekly mean for any hour of the week, together with 10-90 percentiles. In the estimation, we assume the latent states of the additive and multiplicative models approximate the weekly average spot price and log spot price, respectively, and we are left to estimate α_j $j \in \mathcal{J}$. We then can do this by regressing hourly deviations from the weekly spot price mean on the hour of the week (additive), or hourly log deviations from the log mean (multiplicative). We assess the fit by reporting the root mean square error (RMSE) between the point prediction of hourly prices for each model and the actual hourly prices.

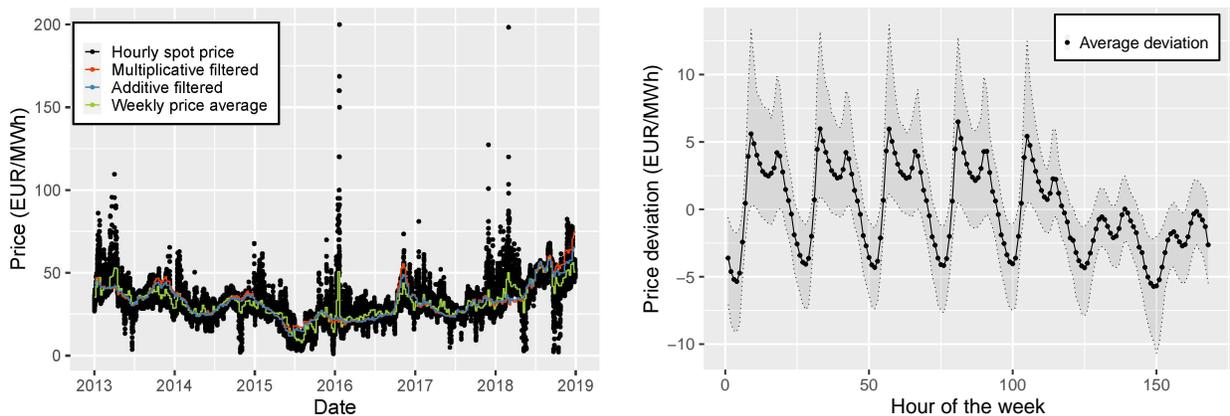
Table 5 reports the RMSE for the hourly spot price prediction from the regression model with time-of-the-week effects in the first column. The additive model performs slightly better. In the last three columns, we have reported the RMSE of model predictions and futures data. Each column represents the average RMSE of short-term contracts (maturity in 1-12 months),

medium-term contracts (maturity in 2-3 years), and long-term contracts (maturity in 4-5 years). For this particular data set, the additive model provides a better fit to short-term contracts and the long-term multiplicative model provides the better fit to long-term contracts. However, we emphasize that the likelihood function is non-convex and sensitive to starting values, and using slightly different data or starting values may lead to substantially different RMSE for both models. We thus conclude that the models perform roughly at par.



(a) Filtered time series after calibration of the additive model, together with the weekly average spot price. (b) Filtered time series after calibration of the multiplicative model, together with the weekly average spot price.

Figure 14: Filtered time series.



(a) Hourly spot price together with estimated latent states and historical weekly price average. (b) Average deviation from weekly average, and 10th and 90th percentiles.

Figure 15: Hourly spot prices.

D. Numerical assessment of Assumption 1

We numerically assess the potential approximation errors by setting the initial states of short-term factors to their long-term means. For simplicity, we normalize the production quantity, which means that the price and the revenue generated by the investment are equivalent. For large ν , the

Table 5: Root mean square error. Data for the day-ahead market include hourly spot prices in the period from 2013-2018. Data for the financial market include contracts with maturities 1, 2, 3, 4, 6, 8, 10, 12, 16, 20, 24, 30, 36, 48, and 60 months from synthetic forward curves, traded weekly in the same period as spot prices.

Model	Day-ahead	Financial market		
	Hourly prices	Short-term contracts	Medium-term contracts	Long-term contracts
Multiplicative	4.29	1.70	1.75	0.99
Additive	4.24	1.61	1.58	1.73

expression for futures prices in (39)-(41) approximates to

$$\hat{F}_{0,\nu} = \exp \left(\xi_0 + \mu_\xi^* \nu - \frac{\lambda_\chi}{\kappa_\chi} + \frac{\sigma_\chi^2}{2\kappa_\chi} + \frac{\rho_{\chi\xi} \sigma_\chi \sigma_\xi}{\kappa_\chi} \right), \quad (54)$$

In order for the discounted price to be bounded as k grows, we need a risk free rate of return $r > \mu_\xi - \lambda_\xi$. Otherwise, the investment will never occur. Figure 16 shows the percentage over- and under estimation of the project value by assuming that investment decisions are unaffected by short-term price deviations. The underestimation stems from not implementing a project that would have generated higher revenue during the first years of the infinite project lifetime, because of a high short-term state. The overestimation stems from possibly implementing a project that will generate lower revenue during the first years because the short-term price is lower than the equilibrium price level. Therefore, to assess the approximation error of Assumption 1 we calculate

$$\frac{\int_0^M e^{-r\nu} (F_{0,\nu} - \hat{F}_{0,\nu}) d\nu}{\int_0^M e^{-r\nu} F_{0,\nu} d\nu}, \quad (55)$$

using a risk-free rate r slightly higher than the risk-adjusted long-term drift. We use $M = 1000$ as a proxy for the infinite lifetime. We have plotted the percentage difference in the interval $[-1, 1]$ for the short-term price factor in Figure 16. By looking at the historical estimates of the latent short-term price factor in Figure 14b, the interval $[-1, 1]$ contains the short-term states in the period from 2013-2018. From Figure 16 we observe that the project value gets overestimated by 0.2% and underestimated by 0.45% for maximum and minimum price deviations in the interval $[-1, 1]$, respectively.

Thus, we find Assumption 1 reasonable with respect to the short-term price deviations. Similarly, we argue that this also is a reasonable assumption for inflow deviations from the mean, since the inflows in our case study possess smaller serial correlations than short-term price deviations. This can be seen by comparing the mean reversion coefficients κ_χ and κ_ζ in Table 4. Therefore, we exclude the level of exogenous short-term factors in the capacity upgrade formulation and set

$$\lim_{t \rightarrow \infty} \mathbb{E}(\chi_t | \chi_0) = 0, \quad \lim_{t \rightarrow \infty} \mathbb{E}(\zeta_t | \zeta_0) = 0, \quad (56)$$

as initial conditions for the short-term exogenous factors for resource valuation to estimate the annual revenue for a representative year.

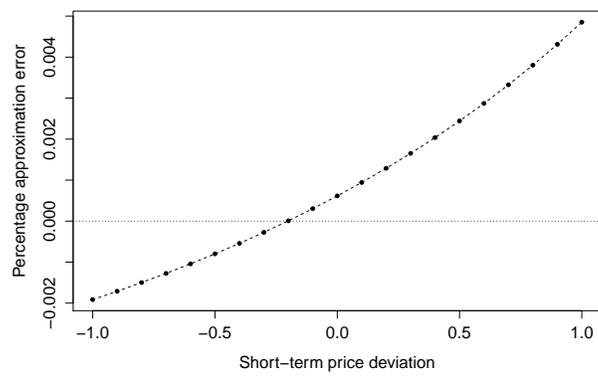


Figure 16: Potential approximating error stemming from Assumption 1