

A Planner-Trader Decomposition for Multi-Market Hydro Scheduling

Kilian Schindler

Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, Switzerland
kilian.schindler@epfl.ch

Napat Rujeerapaiboon

Department of Industrial Systems Engineering and Management, National University of Singapore, Singapore
napat.rujeerapaiboon@nus.edu.sg

Daniel Kuhn

Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, Switzerland
daniel.kuhn@epfl.ch

Wolfram Wiesemann

Imperial College Business School, Imperial College London, United Kingdom
ww@imperial.ac.uk

Peak/off-peak spreads on European electricity forward and spot markets are eroding due to the ongoing nuclear phaseout in Germany and the steady growth in photovoltaic capacity. The reduced profitability of peak/off-peak arbitrage forces hydropower producers to recover part of their original profitability on the reserve markets. We propose a bi-layer stochastic programming framework for the optimal operation of a fleet of interconnected hydropower plants that sells energy on both the spot and the reserve markets. The outer layer (the *planner's problem*) optimizes end-of-day reservoir filling levels over one year, whereas the inner layer (the *trader's problem*) selects optimal hourly market bids within each day. Using an information restriction whereby the planner prescribes the end-of-day reservoir targets one day in advance, we prove that the trader's problem simplifies from an infinite-dimensional stochastic program with 25 stages to a finite two-stage stochastic program with only two scenarios. Substituting this reformulation back into the outer layer and approximating the reservoir targets by affine decision rules allows us to simplify the planner's problem from an infinite-dimensional stochastic program with 365 stages to a two-stage stochastic program that can conveniently be solved via the sample average approximation. Numerical experiments based on a cascade in the Salzburg region of Austria demonstrate the effectiveness of the suggested framework.

Key words: Hydro Scheduling; Reserve Markets; Planner-Trader Decomposition; Stochastic Programming.

1. Introduction

For many decades, hydropower plants have offered a flexible source of emission-free electricity at low running costs. With a world-wide output of 4,246.4 TWh, hydropower accounted for 60.4% of all renewable generation in 2019 (IEA 2020a), and several new hydropower plants are currently being

taken into service in China, Lao People’s Democratic Republic, Portugal and Turkey (IEA 2020b). Pumped-storage plants, which can flexibly transfer water both downstream (through generators, thereby producing energy) and upstream (through pumps, thus consuming energy), account for approximately 95% of the world-wide grid energy storage (DOE 2020).

Traditionally, pumped-storage hydropower plants benefit from releasing the water downstream for electricity generation at peak times and by pumping the water upstream during off-peak periods for future generation (thus adopting a ‘*buy low and sell high*’ strategy). In doing so, the generation companies can exploit the gaps between peak and off-peak electricity prices to make immediate profits. However, these price spreads have been eroding in several European markets since 2008 for two reasons (Mayer 2014): (i) the phaseout of nuclear power plants in Germany has increased base load electricity prices, and (ii) the rapid growth in photovoltaic capacity, whose supply significantly overlaps with the daytime peak hours on weekdays, has reduced the peak electricity prices (Morris and Pehnt 2015, Wirth 2016). As a result, hydropower producers have increasingly turned towards the reserve markets to maintain their profitability.

Electricity is traded at multiple time frames. Besides the forward markets, which trade electricity weeks or months ahead of time, the majority of electricity supply and demand is settled day-ahead on *spot markets*. In practice, however, this cannot be achieved without forecasting errors: unpredicted changes in residential and industrial electricity consumption, failures of power plants, transformers and transmission lines as well as the intermittency of renewable energy sources (such as wind and solar power) all imply that market participants may at times be unable to implement their original plans. These errors are resolved on intra-day markets, which are similar to the spot market but operate with shorter lead times, as well as the *reserve markets*. While the details differ slightly by country, reserve markets are usually divided into primary (frequency containment reserve), secondary and tertiary (both frequency restoration) reserves. Both the primary reserve, which is typically activated within 30 seconds, and the secondary reserve, which is typically activated within 5 minutes, consider time spans of about 15 minutes and are activated automatically. The tertiary reserve, on the other hand, is activated manually within 15 minutes, and it considers time spans of up to one (or, in exceptional cases, several) hour(s). Primary reserves are used to swiftly restore minor grid frequency deviations. Larger deviations are managed via the secondary reserve, and—in the case of major, persistent disruptions and only when the secondary reserves fail to restore the balance between demand and supply—the tertiary reserve. We refer to Aasgård et al. (2019), European Commission (2020) and Fernández-Muñoz et al. (2020) for further details on the different energy markets.

Generation companies benefit from trading in the reserve markets in two ways. First, depending on the market, they may receive capacity fees regardless of whether the reserve capacities are

activated. Second, they earn money proportional to the increase or decrease in their production levels when the reserves are activated. Although most thermal plants can provide all three types of reserve when they are online, their ability to do so is limited by their current power output (which must be below capacity to offer upward corrections and above the minimum stable operation limit to offer downward corrections in the energy output) as well as their ramp rate. Gas turbines are the only thermal plants that can offer tertiary reserve when they are offline. With a start-up time of a few minutes and high ramp rates, on the other hand, hydropower plants are an ideal source for any type of reserve, and they can provide secondary and tertiary reserve even when they are offline. Participation on the reserve markets should ease the pressure on the hydropower plant operators, who are struggling to recover their capital costs as well as their original profitability on the forward and spot markets because of the eroding peak/off-peak spreads.

In this paper, we develop a stochastic program that maximizes the revenues earned by a hydropower producer from simultaneously trading in both the spot and the reserve markets. The resulting optimization problem is computationally challenging for at least three reasons. First, the problem involves a large number of decision stages. Indeed, a planning horizon of one year is required to account for the seasonality of rainfall and electricity prices, while generation and pumping decisions need to be taken on an hourly basis. Second, the problem is affected by significant uncertainty in the electricity prices, the timing and amount of the water inflows (*i.e.*, rainfall and snowmelt) as well as the reserve activations. Finally, the problem has random recourse since the water flows between the reservoirs not only depend on the bidding decisions in the spot and reserve markets, but they also depend on the reserve activations by the transmission system operator.

To obtain a tractable solution scheme, we equivalently express the stochastic program of the hydropower producer as a bi-layer problem where the outer layer *planner's problem* is a 365-stage stochastic program with daily resolution that determines end-of-day water levels for all reservoirs, and the inner layer *trader's problem* is a 25-stage stochastic program with hourly resolution that uses the target water levels to determine optimal bids for the spot and reserve markets. Using an information restriction whereby the planner prescribes the target water levels one day in advance, we prove that the trader's problem can equivalently be reduced to a two-stage stochastic program with only two scenarios. By substituting this reduced trader's problem back into the planner's problem, applying an exact perfect information relaxation to all hourly bidding decisions and conservatively approximating the daily water level targets via linear decision rules, we obtain a two-stage stochastic program that can be solved efficiently via the sample average approximation.

The main contributions of this paper may be summarized as follows.

- (i) We propose a stochastic program that maximizes the revenues of a hydropower producer who simultaneously trades in the spot and reserve markets and whose fleet of plants can have an

arbitrary topology. Our model faithfully accounts for uncertainty in the electricity prices, the water inflows as well as the reserve activations. To our best knowledge, we propose the first model that accounts for the stochasticity of the reserve activations at an hourly granularity.

- (ii) Under the assumption that the end-of-day water levels are chosen one day in advance, we develop a novel planner-trader decomposition that allows us to separately study the reservoir management and trading problems. We equivalently reformulate the trader’s problem as a tractable linear program, and we approximate the planner’s problem through a combination of sample average approximation and linear decision rules. The resulting model can be readily solved within a few minutes on standard hardware, which is in stark contrast to the vast majority of stochastic hydropower models that account for reserve market participation.
- (iii) We report initial numerical results for a cascade of three connected reservoirs operating in the control area of the Austrian Power Grid AG. Our results suggest that significant revenues can be earned on the reserve markets if the stochastic reserve activations are modeled faithfully.

The detailed modeling of uncertainty, which is essential to faithfully model the stochasticity of reserve activations, requires us to make a number of simplifying assumptions. First, we assume that the hydropower producer only participates in the spot and the secondary reserve market. While we envision that the inclusion of the intra-day market and the primary and tertiary reserve markets is principally possible, they would result in larger optimization models that likely demand substantially longer solution times. The inclusion of a forward market, on the other hand, appears to be fundamentally more difficult as it would require the consideration of different time scales. Second, we simplify the bidding process in both the spot and the reserve markets. Instead of directly modeling the possibility to bid at multiple price levels, we study a single-bid strategy that is finetuned out-of-sample by two hyperparameters. Third, we omit a number of technical aspects, such as the discrete (on/off) nature of equipment operations, head effects and water flow delays. Such considerations often lead to the inclusion of integer variables and/or non-linear relationships between variables, which would significantly complicate the solution of the stochastic program.

The liberalization of energy markets world-wide has sparked significant academic interest in the optimal management of hydropower plants. Early works focus on participation in the spot markets, and they include scenario tree approaches (Andrieu et al. 2010, Carpentier et al. 2013, Aasgård et al. 2014), stochastic dynamic programming (Pritchard et al. 2005) and stochastic dual dynamic programming approaches (Pereira and Pinto 1991, Gjelsvik et al. 2010, Philpott and de Matos 2012, Shapiro et al. 2013, de Matos et al. 2015, Philpott et al. 2018) as well as robust and stochastic programs using affine decision rules (Pan et al. 2015, Lorca et al. 2016, Gauvin et al. 2017, 2018).

In departure from the aforementioned single-market studies, Löhndorf et al. (2013) consider a fleet of hydropower plants that participates in both the spot and the intra-day market. The problem

comprises uncertain water inflows and electricity prices. The time horizon of one year is subdivided into inter-day reservoir management problems (at a daily resolution) as well as intra-day bidding problems (at an hourly resolution). The authors solve the problem by a combination of stochastic dual dynamic programming and approximate dynamic programming. The runtimes are reported to range between 4 and 40 hours.

The recent shift of the hydropower producers' attention towards the reserve markets has led to a number of multi-market models that consider both the spot and the reserve markets. Below, we distinguish between deterministic models and stochastic models using scenario trees, stochastic dynamic programming as well as stochastic dual dynamic programming.

In return for disregarding the stochasticity of the problem data, deterministic hydropower models typically offer a superior representation of the physical details of a hydropower plant, such as the consideration of discrete operation modes, head effects as well as a distinction between fixed-speed and variable-speed units. Chazarra et al. (2014, 2018) propose daily mixed-integer programs for pumped-storage hydropower plants in Spain that model the hourly bidding on the spot and secondary reserve markets. Schillinger et al. (2017) apply a yearly linear programming reservoir management problem with weekly nonlinear bidding subproblems that comprise the spot, intra-day as well as primary, secondary and tertiary reserve markets to a hydropower plant in Switzerland. Fodstad et al. (2018) propose a deterministic linear program with a one year time horizon that determines the optimal hourly bidding on the day-ahead, intra-day and reserve markets for a hydropower plant in Norway. Since deterministic models neglect the uncertainty in the water inflows, electricity prices and reserve activations, they adopt a perfect information relaxation that may lead to overly optimistic estimates of the revenues achievable by a hydropower plant.

Chazarra et al. (2016) develop a two-stage stochastic program for the short-term scheduling of a fleet of hydropower plants that participates in the spot and the reserve markets. The model comprises a time horizon of one week, where the decisions of the first day are taken in the first stage and the decisions of the remaining six days are taken in the second stage, respectively. Chazarra et al. (2018) propose a single-stage stochastic program for a hydropower plant that participates in the spot and secondary reserve market. The authors subdivide the planning horizon into daily problems with hourly resolution. The daily problems optimize the conditional value-at-risk of the revenues over 100 scenarios, subject to uncertainty in the spot and reserve prices. The runtime of the daily model varies between 3 and 17 minutes. Klæboe et al. (2019), finally, consider a three-stage stochastic program where the hydropower producer submits the spot and reserve bids in the first and third stage, respectively, while the spot market prices and reserve market prices are observed in the second and third stage, respectively. The authors report computation times between 2 and 7 hours. While scenario-tree based stochastic programs allow to capture complex

relationships between the uncertain problem parameters, the scenario discretization typically fails to provide an implementable policy, and the problem size tends to grow exponentially with the considered time horizon. As a result, most studies consider a small number of stages and model the stochastic price and reserve activation processes at a coarser (*e.g.*, daily) granularity.

Abgottspon and Andersson (2012) propose a stochastic dynamic program for a hydropower plant with a single reservoir that participates in the spot and secondary reserve markets. The authors consider a yearly problem whose weekly time stages are connected through a single state variable describing the water level of the reservoir. Each weekly problem is modelled as a two-stage stochastic program whose first stage decides upon the reserve market offering under uncertain prices and whose second stage manages the water flows. The authors report that the entire model can be solved within 20 minutes via large-scale parallelization. Similar to Abgottspon and Andersson (2012), Helseth et al. (2017) study a hydropower plant with a single reservoir that participates in the spot and secondary reserve markets. The authors also consider a yearly problem with weekly time periods, however the state variables now comprise the weekly average energy prices in addition to the reservoir level. The authors report computation times of 24 to 40 hours. While stochastic dynamic programs allow to model complex non-convex phenomena in the weekly subproblems, the curse of dimensionality implies that the runtimes increase quickly with the dimension of the state space, which makes the stochastic dynamic programs primarily suited for small fleets of plants with a coarse description of the involved stochastic processes. Moreover, the need to discretize the water levels can introduce numerical inaccuracies.

Abgottspon et al. (2014) compare the performance of a stochastic dynamic and a stochastic dual dynamic program that both manage a hydropower plant with a seasonal and a daily reservoir over a yearly time horizon with weekly time periods. The weekly subproblems are modeled as two-stage stochastic programs that decide upon the spot and reserve market participation in an hourly resolution under uncertain water inflows and energy prices. The overall models record the water level of the seasonal reservoir in the state variable and optimize a nested conditional value-at-risk. Computation times are not reported. Helseth et al. (2015) propose a stochastic dual dynamic program that manages a multi-reservoir fleet of hydro plants over the course of one year in weekly time stages. The state variables of the problem record the weekly water levels of the reservoirs as well as the spot prices, while the weekly subproblems manage the commitments on the spot and the primary reserve markets. The overall problem is solved in 42 hours. In a follow-up paper, Helseth et al. (2016) modify the stochastic dual dynamic program of Helseth et al. (2015) to incorporate stochastic reserve prices. The resulting model can be solved within 55-60 hours. While stochastic dual dynamic programs scale better than stochastic dynamic programs, they too tend to be limited to small and medium-sized fleets that are managed over coarse (typically weekly)

time intervals with relatively crude approximations of the involved stochastic processes. This is particularly critical for the reserve market activations in small (*e.g.*, daily) reservoirs, where a daily or weekly time granularity may be too coarse to ensure an implementable (hourly) policy.

The remainder of the paper unfolds as follows. Section 2 models the overall bidding problem of the hydropower producer. Section 3 transforms this formulation to an equivalent bi-layer stochastic program. Section 4 reformulates the inner layer of this stochastic program as a tractable linear program. Section 5 presents a tractable approximation for the outer layer of the stochastic program as well as a rolling horizon solution framework. Section 6 applies our method to a hydropower cascade in Austria. Auxiliary results and all proofs, finally, are relegated to the appendix.

Notation. We denote by $\mathbf{0}$ and $\mathbf{1}$ the appropriately sized vectors of all zeros and all ones, respectively. The Hadamard product is denoted by “ \circ ” and refers to the element-wise vector multiplication. For any $x \in \mathbb{R}$, we set $x^+ = \max\{x, 0\}$ and $x^- = \max\{-x, 0\}$ such that $x = x^+ - x^-$. All random objects are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ consisting of a sample space Ω , a σ -algebra $\mathcal{F} \subseteq 2^\Omega$ of events and a probability measure \mathbb{P} on \mathcal{F} . For a σ -algebra $\mathcal{G} \subseteq \mathcal{F}$, we denote by $\mathcal{L}^k(\mathcal{G})$ the set of all integrable, \mathcal{G} -measurable functions $g: \Omega \rightarrow \mathbb{R}^k$, and if $k = 1$, we simply write $\mathcal{L}(\mathcal{G})$.

2. Bidding Model

We consider a hydropower generation company that operates a cascade of reservoirs over a time horizon spanning at least one year to account for the seasonality of electricity prices and water inflows. We partition the planning horizon into days indexed by $d \in \mathcal{D} := \{1, \dots, D\}$, as well as into hours indexed by $t \in \mathcal{T} := \{1, \dots, T\}$. To simplify the exposition, we assume that T is divisible by D , that is, the planning horizon accommodates an integral number of days. We further denote by $H = T/D$ the number of hours per day, but we emphasize that all subsequent results remain valid for any $H \in \mathbb{N}$, *e.g.*, if we partition each day into 15 minute intervals. The present time is modeled by a fictitious hour $t = 0$ on day $d = 0$. It is followed by hour $t = 1$ on day $d = 1$, that is, the first hour of the planning horizon. No decisions are taken at time $t = 0$; its inclusion merely serves to unify our exposition so that decisions and constraints at time $t = 1$ can be expressed similar to those in any other time period. For any $t, t' \in \{0\} \cup \mathcal{T}$ with $t \leq t'$, we denote the set of hours between and including t and t' as $[t, t'] = \{t, t+1, \dots, t'\}$. Note in particular that $[t, t] = \{t\}$. Furthermore, we denote by $d(t)$ the day containing hour t and by $\mathcal{T}(d)$ the set of all hours in day d . Finally, we introduce the following functions of $d \in \mathcal{D}$ and $t \in \mathcal{T}$:

$$\begin{aligned} \uparrow(d) &= \text{last hour of day } d, & \uparrow(t) &= \text{last hour of day } d(t), \\ \downarrow(d) &= \text{first hour of day } d, & \downarrow(t) &= \text{first hour of day } d(t), \\ \Downarrow(d) &= \text{last hour of day } d-1, & \Downarrow(t) &= \text{last hour of day } d(t)-1. \end{aligned}$$

By slight abuse of notation, the operators $\uparrow(\cdot)$, $\downarrow(\cdot)$ and $\Downarrow(\cdot)$ are defined both on \mathcal{D} and \mathcal{T} . The correct interpretation of these operators will always be clear from the context. Observe also that the two $\Downarrow(\cdot)$ operators are not strictly needed because they can be expressed in terms of $\downarrow(\cdot)$ via the identities $\Downarrow(d) = \uparrow(d-1)$ and $\Downarrow(t) = \uparrow(d(t)-1)$. However, they are useful to simplify notation.

We represent the topology of the interconnected reservoirs by a directed acyclic graph with a set of nodes \mathcal{R} representing the reservoirs and a set of arcs $\mathcal{A} \subseteq \mathcal{R} \times \mathcal{R}$ representing the hydraulic connections between the reservoirs. We denote the cardinalities of \mathcal{R} and \mathcal{A} by R and A , respectively. A tuple (r, r') is an arc in \mathcal{A} if r is an upstream reservoir of r' (and hence r' is a downstream reservoir of r). Without loss of generality, we assume that each arc is equipped with a generator that converts kinetic energy of water flowing downstream into electric energy and with a pump that uses electric energy to lift water upstream. Note that some of these generators and pumps may have zero capacity, which models situations in which these devices are absent. Without loss of generality, we further assume that the node R is the unique sink of the graph. This dummy reservoir has infinite volume and models the discharge into a river at the end of the reservoir cascade. Accordingly, arcs discharging into this dummy reservoir may be equipped with generators, but not with pumps. In line with real hydropower plants, we assume that every reservoir (except the aforementioned dummy reservoir) has at least one outgoing arc along which water can be discharged whenever required. The topology of the cascade can be encoded conveniently through an $R \times A$ incidence matrix \mathbf{M} , where for each $(r, a) \in \mathcal{R} \times \mathcal{A}$ we have that

$$[\mathbf{M}]_{r,a} = m_{r,a} = \begin{cases} -1 & \text{if arc } a \text{ leaves reservoir } r, \\ +1 & \text{if arc } a \text{ enters reservoir } r, \\ 0 & \text{otherwise.} \end{cases}$$

Each reservoir is characterized by—potentially time-dependent—lower and upper bounds on its filling level. The time dependence of these bounds may reflect changing safety margins that are imposed to account for the seasonality of inflows, scheduled maintenance work or environmental regulations. For each hour $t \in \mathcal{T}$, we collect the lower and upper reservoir bounds in the vectors $\underline{w}_t \in \mathbb{R}_+^R$ [m³] and $\overline{w}_t \in (\mathbb{R}_+ \cup \{+\infty\})^R$ [m³], respectively. The bounds for the dummy reservoir are $\underline{w}_{t,R} = 0$ and $\overline{w}_{t,R} = +\infty$, that is, the river can absorb arbitrary amounts of water. The generators installed along the arcs of our cascade are characterized by the upper bounds $\overline{g}_t \in \mathbb{R}_+^A$ [m³] on the hourly water flows and the generator efficiencies $\boldsymbol{\eta}_t \in \mathbb{R}_+^A$ [MWh/m³]. Similarly, the pumps are described by the upper bounds $\overline{p}_t \in \mathbb{R}_+^A$ [m³] on the hourly water flows and the *inverse* pump efficiencies $\boldsymbol{\zeta}_t \in \mathbb{R}_+^A$ [MWh/m³]. The laws of thermodynamics dictate that $\boldsymbol{\eta}_t < \boldsymbol{\zeta}_t$ for all $t \in \mathcal{T}$.

The decision problem of the hydropower generation company is affected by three types of risk factors with finite expectations, namely fluctuating market prices, uncertain reserve allocations

and reserve activations as well as natural inflows into the reservoirs caused by unpredictable meteorological phenomena. Below, we first describe the uncertainties related to the different markets.

We assume that the hydropower producer participates in the spot market and in the secondary reserve markets for up- and down-regulation. In central Europe, allocations and prices on these markets are determined by separate auctions for the different hours of the day (consecutive hours could be merged to blocks). The spot market is mediated via a sealed-bid pay-as-clear auction, where demand and supply bids correspond to tuples comprising an energy quantity to be sold or bought and a price. Bids are accepted based on the spot price, which is found by matching demand and supply. The secondary reserve markets, on the other hand, are mediated via open pay-as-bid auctions, where bids correspond to tuples comprising an energy quantity offered as reserve capacity, a capacity price and an activation price. Bids are accepted in the order of their capacity prices until the regulation demand is covered. If grid imbalances require an intervention, the accepted bids are activated in the order of their energy prices until the imbalances are eliminated. This means that bids with low energy prices are more likely to be activated than bids with high energy prices.

A faithful representation of the different auction mechanisms would lead to an excessively complicated model. In order to obtain a realistic model that is amenable to numerical solution, we use the following approximations. First, we assume that the spot and reserve markets are infinitely liquid. Second, we model the spot price in any hour t as an exogenous random variable $\pi_t^s \in \mathbb{R}$ [\$/MWh] that is revealed one day in advance at time $\downarrow(t)$, and we assume that the energy quantities traded in hour t on the spot market may adapt to π_t^s . This approximation is reasonable if the producer has no market power and places a continuum of bids corresponding to all possible spot price realizations. Third, we model the capacity prices for up- and down-regulation in any hour t as exogenous random variables $\pi_t^u \in \mathbb{R}$ [\$/MWh] and $\pi_t^v \in \mathbb{R}$ [\$/MWh], respectively, which are revealed one day in advance at time $\downarrow(t)$, and we assume that the energy quantities offered as reserve capacity in hour t may adapt to π_t^u and π_t^v . By slight abuse of terminology, we will sometimes refer to these energy quantities simply as ‘bids.’ We also assume that the allocation of reserve capacity in hour t is random and can be modeled through exogenous Bernoulli random variables κ_t^u and κ_t^v with success probabilities $\hat{\kappa}_t^u = \mathbb{E}[\kappa_t^u]$ and $\hat{\kappa}_t^v = \mathbb{E}[\kappa_t^v]$, respectively, which are revealed at time $\downarrow(t)$. The allocated reserve capacity for up- or down-regulation is then obtained by multiplying the offered reserve capacity with the corresponding Bernoulli random variable. This approximation is reasonable if the producer has no market power and places only one bid on the reserve-up market with capacity price π_t^u and only one bid in the reserve-down market with capacity price π_t^v . In reality, the prices π_t^u and π_t^v constitute decisions that have a diminishing impact on the allocation probabilities $\hat{\kappa}_t^u$ and $\hat{\kappa}_t^v$, respectively. Treating these quantities as exogenous parameters amounts

to prescribing a fixed trade-off between the probability that a given bid is accepted and the financial compensation that is earned if the bid is accepted. In the model to be developed below we will optimize this trade-off via a procedure reminiscent of hyperparameter tuning. As the fourth and last approximation, we model the activation prices for up- and down-regulation as exogenous random variables $\psi_t^u \in \mathbb{R}$ [\$/MWh] and $\psi_t^v \in \mathbb{R}$ [\$/MWh], respectively, which are revealed in real time. In addition, we model the percentages of time in which the system operator activates up- and down-regulation in hour t as continuous random variables ρ_t^u and ρ_t^v with supports $[0, \bar{\rho}_t^u]$ and $[0, \bar{\rho}_t^v]$ and with mean values $\hat{\rho}_t^u = \mathbb{E}[\rho_t^u]$ and $\hat{\rho}_t^v = \mathbb{E}[\rho_t^v]$, respectively. The activations ρ_t^u and ρ_t^v are also revealed in real time. The amount of energy activated for up- and down-regulation in hour t is obtained by multiplying the allocated reserve capacity with ρ_t^u and ρ_t^v , respectively. This approximation is again reasonable if the producer has no market power and places only one bid on the reserve-up market with activation price ψ_t^u and only one bid in the reserve-down market with activation price ψ_t^v . In reality, ψ_t^u and ψ_t^v constitute decisions that have a diminishing impact on the expected activations $\hat{\rho}_t^u$ and $\hat{\rho}_t^v$, respectively. Treating these quantities as exogenous parameters amounts to prescribing a fixed trade-off between the probability that an accepted reserve bid is activated and the financial compensation that is earned if the bid is activated. In the model to be developed below we will optimize this trade-off via a procedure reminiscent of hyperparameter tuning. Note that, as any bidding zone may either face an over-supply or an under-supply of energy but not both at the same time, the sum of ρ_t^u and ρ_t^v cannot exceed 1. Throughout the paper we thus assume that the pairs of reserve activations on the up- and down-regulation markets have support $\{(\rho_t^u, \rho_t^v) \in [0, \bar{\rho}_t^u] \times [0, \bar{\rho}_t^v] : \rho_t^u + \rho_t^v \leq 1\}$, where the upper bounds $\bar{\rho}_t^u$ and $\bar{\rho}_t^v$ can be set, for example, to quantiles of the historical hourly activation percentages. We also assume that the reserve activations are serially independent as well as independent of all other exogenous uncertainties. This assumption may only hold approximately in practice, and we discuss its validity in Appendix A.

Finally, we denote the natural inflows into the various reservoirs by the vector $\phi_t \in \mathbb{R}_+^R$ [m³], and we assume that they are revealed a day in advance at time $\downarrow(t)$ thanks to an accurate forecast. From now on we gather all random variables that are revealed in hour $t \in \{0\} \cup \mathcal{T}$ in the vector

$$\xi_t = \begin{cases} (\{\pi_\tau^s, \pi_\tau^u, \pi_\tau^v, \phi_\tau\}_{\tau=1}^H) & \text{if } t = 0, \\ (\psi_t^u, \psi_t^v, \rho_t^u, \rho_t^v, \{\pi_\tau^s, \pi_\tau^u, \pi_\tau^v, \phi_\tau\}_{\tau=t+1}^{t+H}) & \text{if } t \in \left[\bigcup_{\tau \in \mathcal{T}} \downarrow(\tau) \right] \setminus \{0\}, \\ (\psi_t^u, \psi_t^v, \rho_t^u, \rho_t^v, \{\kappa_\tau^u, \kappa_\tau^v\}_{\tau=t}^{t+H-1}) & \text{if } t \in \left[\bigcup_{\tau \in \mathcal{T}} \downarrow(\tau) \right], \\ (\psi_t^u, \psi_t^v, \rho_t^u, \rho_t^v) & \text{otherwise.} \end{cases}$$

In other words, the spot and capacity prices as well as the inflows of day $d \in \mathcal{D}$ are revealed in the last hour of the previous day, the reserve allocations of day d are revealed in the first hour of the

current day, whereas the activation prices and actual activations are revealed hourly in real time. We gather all random variables whose values are revealed within the interval $[t, t']$ in the vector $\boldsymbol{\xi}_{[t, t']} = (\boldsymbol{\xi}_t, \dots, \boldsymbol{\xi}_{t'})$. Furthermore, we capture the information revealed during this time interval by the σ -algebra $\mathcal{F}_{[t, t']} = \sigma(\boldsymbol{\xi}_{[t, t']})$, *i.e.*, the σ -algebra generated by $\boldsymbol{\xi}_{[t, t']}$. For ease of notation, we abbreviate $\boldsymbol{\xi}_{[0, t]}$ and $\mathcal{F}_{[0, t]}$ by $\boldsymbol{\xi}_{[t]}$ and $\mathcal{F}_{[t]}$, respectively.

The hydropower generation company must decide each day how much energy to offer on the different markets in every hour of the following day. This in turn necessitates a strategy for operating the reservoir system that is guaranteed to honor all market commitments. We denote the bids placed on the spot market, the reserve-up market and the reserve-down market by $s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[\psi(t)]})$ [MWh], respectively. Our model involves the water flows along the arcs as well as the reservoir filling levels as real-time operational decisions. Specifically, we denote the generation flows, *i.e.*, the amounts of water running through the turbines in each hour, by $\mathbf{g}_t \in \mathcal{L}^A(\mathcal{F}_{[t]})$ [m³] and the consumption flows, *i.e.*, the amounts of water running through the pumps in each hour, by $\mathbf{p}_t \in \mathcal{L}^A(\mathcal{F}_{[t]})$ [m³]. Furthermore, we ignore concerns of flooding related to the dummy reservoir R and assume that one can spill unlimited amounts of water $\mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[t]})$ [m³] along all arcs from the upstream reservoirs to the downstream reservoirs/ivers. The reservoir filling levels at the end of each hour are denoted by $\mathbf{w}_t \in \mathcal{L}^R(\mathcal{F}_{[t]})$ [m³]. The initial filling levels $\mathbf{w}_0 \in \mathbb{R}^R$ [m³] are deterministic input parameters of our model. We assume that the hydropower generation company is risk-neutral but aims to satisfy all constraints robustly—especially those related to reserve market commitments. This appears justified as non-compliance would incur high penalty fees or even lead to an exclusion from the reserve markets. Thus, the company maximizes the expected cumulative revenues across the entire planning horizon while maintaining its ability to respond to all possible call-offs on the reserve markets. The objective function thus takes the form

$$\sum_{t \in \mathcal{T}} \mathbb{E} \left[\pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t \right],$$

and the non-anticipativity constraints to be respected are

$$s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[\psi(t)]}), \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[t]}), \mathbf{w}_t \in \mathcal{L}^R(\mathcal{F}_{[t]}) \quad \forall t \in \mathcal{T}.$$

While spot market bids have no sign restrictions (*i.e.*, the company can both buy or sell electricity in the spot market), bids in the reserve markets as well as the operational flow decisions must be non-negative. Also, the flow decisions must obey the capacity limits of the respective generators and pumps, and the reservoir levels must stay within the respective bounds. Thus, we require

$$0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t, \underline{\mathbf{w}}_t \leq \mathbf{w}_t \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.}$$

In addition, the hydropower company is obliged to produce and/or consume energy according to its market commitments. If the same system operator is responsible for both the spot and the reserve markets, then the market commitments are enforced through the following constraint

$$s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.}$$

Taking into account the natural inflows, the reservoir filling levels obey the dynamic equation

$$\mathbf{w}_t = \mathbf{w}_{t-1} + \boldsymbol{\phi}_t + \mathbf{M}(\mathbf{g}_t - \mathbf{p}_t + \mathbf{z}_t) \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.}$$

In summary, the company aims to solve the stochastic optimization problem

$$\begin{aligned} & \sup \sum_{t \in \mathcal{T}} \mathbb{E} [\pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t] \\ & \text{s.t. } s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[\downarrow(t)]}), \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[t]}), \mathbf{w}_t \in \mathcal{L}^R(\mathcal{F}_{[t]}) \quad \forall t \in \mathcal{T} \\ & 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.} \\ & s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.} \\ & \mathbf{w}_t = \mathbf{w}_{t-1} + \boldsymbol{\phi}_t + \mathbf{M}(\mathbf{g}_t - \mathbf{p}_t + \mathbf{z}_t) \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.} \\ & \underline{\mathbf{w}}_t \leq \mathbf{w}_t \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}, \mathbb{P}\text{-a.s.} \end{aligned} \tag{C}$$

3. Planner-Trader Decomposition

In this section we will demonstrate that problem (C) can be decomposed into subproblems that lend themselves to further simplification. Specifically, we first express (C) as a bi-layer stochastic program, where the outer layer (the *planner's problem*) optimizes over the reservoir filling levels with a daily granularity, whereas the inner layer (the *trader's problem*) optimizes over the bidding decisions with an hourly granularity. To this end, recall that the decision variables $\{\mathbf{w}_t\}_{t \in \mathcal{T}}$ in problem (C) represent the end-of-hour filling levels of the reservoirs. In order to reformulate (C), we will overload notation and denote by $\mathbf{w}_d = \mathbf{w}_{\uparrow(d)}$, $\underline{\mathbf{w}}_d = \underline{\mathbf{w}}_{\uparrow(d)}$ and $\bar{\mathbf{w}}_d = \bar{\mathbf{w}}_{\uparrow(d)}$ the end-of-day reservoir filling levels and their lower and upper bounds for day $d \in \mathcal{D}$, respectively. Next, we use the reservoir balance constraints to re-express all end-of-hour filling levels $\{\mathbf{w}_t\}_{t \in \mathcal{T}}$ in terms of the end-of-day filling levels $\{\mathbf{w}_d\}_{d \in \mathcal{D}}$, the hourly flow decisions $\{\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t\}_{t \in \mathcal{T}}$ and the natural inflows $\{\boldsymbol{\phi}_t\}_{t \in \mathcal{T}}$:

$$\mathbf{w}_t = \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d).$$

Using this relation, the bidding model (C) can be recast equivalently as

$$\begin{aligned}
& \sup \sum_{t \in \mathcal{T}} \mathbb{E} [\pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t] \\
& \text{s.t. } s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{\downarrow(t)}), \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[t]}), \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{\uparrow(d)}) \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d) \\
& 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\
& s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \quad (1) \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\
& \mathbf{w}_d = \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.} \\
& \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.},
\end{aligned}$$

Note that the end-of-day reservoir bounds in the last line are implied by the end-of-hour reservoir bounds two lines above and are thus redundant. However, they will be instrumental for proving the tightness of a relaxation of problem (1) to be derived below. Next, we introduce another auxiliary problem that imposes the daily reservoir balance constraints as inequalities rather than equalities.

$$\begin{aligned}
& \sup \sum_{t \in \mathcal{T}} \mathbb{E} [\pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t] \\
& \text{s.t. } s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{\downarrow(t)}), \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[t]}), \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{\uparrow(d)}) \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d) \\
& 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\
& s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \quad (2) \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\
& \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.} \\
& \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.}
\end{aligned}$$

The next proposition asserts that the optimization problems (1) and (2) are indeed equivalent. This proposition relies on the observation that any slacks between the left- and right-hand sides of the daily reservoir balance constraints can systematically be eliminated by suitable spillages, which are unrestricted in magnitude by our earlier assumptions.

PROPOSITION 1. *The optimal values of problems (1) and (2) are equal.*

We are now ready to decompose the bidding problem (C), which jointly maximizes over all hourly and daily decisions, into planning and trading subproblems that maximize only over the daily and the hourly decisions, respectively. The proposed decomposition assumes that a fictitious

planner and a fictitious trader collaborate to solve problem (C) in the following manner. The planner determines the end-of-day reservoir levels by solving the stochastic program

$$\begin{aligned} & \sup \sum_{d \in \mathcal{D}} \mathbb{E}[\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})] \\ & \text{s.t. } \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{[\uparrow(d)]}) \quad \forall d \in \mathcal{D} \\ & \quad \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.}, \end{aligned} \tag{CP}$$

which maximizes the expected daily profits accrued over the entire planning horizon. Here, the function $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})$ represents the expected profit earned by the trader on day d , conditional on the information $\boldsymbol{\xi}_{[\downarrow(d)]}$ available at the beginning of the day and conditional on the initial and terminal reservoir levels \mathbf{w}_{d-1} and \mathbf{w}_d , respectively, imposed by the planner. The function $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})$ evaluates the optimal value of the stochastic program

$$\begin{aligned} & \sup \sum_{t \in \mathcal{T}(d)} \mathbb{E}[\pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t \mid \boldsymbol{\xi}_{[\downarrow(d)]}] \\ & \text{s.t. } s_t, u_t, v_t \in \mathbb{R}, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]}) \quad \forall t \in \mathcal{T}(d) \\ & \quad 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\ & \quad s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \tag{CT} \\ & \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\ & \quad \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \end{aligned}$$

solved by the trader, where $\mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}$ denotes the probability measure \mathbb{P} conditioned on $\boldsymbol{\xi}_{[\downarrow(d)]}$, and $\mathbb{E}[\cdot | \boldsymbol{\xi}_{[\downarrow(d)]}]$ denotes the expectation with respect to $\mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}$. Note that the trader's problem (CT) may be infeasible for some reservoir targets \mathbf{w}_{d-1} and \mathbf{w}_d , in which case $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})$ evaluates to $-\infty$ and thus introduces implicit constraints in the planner's problem (CP). We also highlight that problem (CT) constitutes a stochastic program with contextual information (Ban and Rudin 2019, Bertsimas and Kallus 2020), where the contextual covariates $\boldsymbol{\xi}_{[\downarrow(d)]}$ impact the (conditional) distribution of the random variables in the objective and the constraints.

The next proposition asserts that the planner-trader decomposition incurs no loss of optimality.

PROPOSITION 2. *The optimal values of problems (2) and (CP) are equal.*

In conclusion, we emphasize that the proposed planner-trader decomposition results in an exact reformulation of the underlying bidding model (C) and that they involve no approximations. This decomposition is useful because the trader's subproblem can be reduced to a tractable linear program under a mild information restriction.

4. Reduction of the Trader's Problem

The trader's problem (CT) constitutes an infinite-dimensional stochastic program because it optimizes over functional flow decisions. Such problems are known to be intractable in general (Dyer and Stougie 2006, Hanasusanto et al. 2016). In this section we will show, however, that the stochastic program (CT) is equivalent to a finite-dimensional linear program if the planner chooses the reservoir filling levels one day ahead of time. Note that selecting \mathbf{w}_d at the end of day $d-1$ rather than the end of day d amounts to an information restriction whereby the planner sacrifices potentially useful information revealed during day d and thus foregoes some of the achievable expected profit. In other words, requiring that the reservoir filling levels be pre-committed a day in advance leads to a conservative approximation of the planning problem (CP). For later reference, we formally state this approximation below.

Approximation 1 (Information restriction) *For every $d \in \mathcal{D}$, the end-of-day reservoir filling level \mathbf{w}_d is restricted to $\mathcal{L}^R(\mathcal{F}_{\uparrow(d-1)}) = \mathcal{L}^R(\mathcal{F}_{\downarrow(d)}) \subseteq \mathcal{L}^R(\mathcal{F}_{\uparrow(d)})$.*

Under Approximation 1, \mathbf{w}_d becomes a measurable function of $\boldsymbol{\xi}_{\downarrow(d)}$, which captures the information available to the trader at the beginning of day d . Similarly, the natural inflows $\{\phi_t\}_{t \in \mathcal{T}(d)}$ of day d were assumed to be predictable at the beginning of the day and can therefore also be expressed as measurable functions of $\boldsymbol{\xi}_{\downarrow(d)}$. Finally, the coefficients

$$\tilde{\pi}_t^s = \mathbb{E}[\pi_t^s \mid \boldsymbol{\xi}_{\downarrow(d)}], \quad \tilde{\pi}_t^u = \mathbb{E}[\kappa_t^u(\pi_t^u + \rho_t^u \psi_t^u) \mid \boldsymbol{\xi}_{\downarrow(d)}] \quad \text{and} \quad \tilde{\pi}_t^v = \mathbb{E}[\kappa_t^v(\pi_t^v + \rho_t^v \psi_t^v) \mid \boldsymbol{\xi}_{\downarrow(d)}] \quad \forall t \in \mathcal{T}(d)$$

of the (here-and-now) bidding decisions $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(d)}$ in the objective function of problem (CT) constitute measurable functions of $\boldsymbol{\xi}_{\downarrow(d)}$ thanks to the properties of conditional expectations. We may thus conclude that, once a particular realization of $\boldsymbol{\xi}_{\downarrow(d)}$ is fixed, the individual trader's problem (CT) reduces to a multistage stochastic program with $H+1$ decision stages. Here, the reserve allocations $\{(\kappa_t^u, \kappa_t^v)\}_{t \in \mathcal{T}(d)}$, which are revealed simultaneously at the beginning of hour 1, and the reserve activations $\{(\rho_t^u, \rho_t^v)\}_{t \in \mathcal{T}(d)}$, which are revealed sequentially at the beginning of the respective hours, are the only exogenous uncertain parameters affecting the constraints, and the objective function is independent of the (wait-and-see) flow decisions $\{(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$.

Approximation 1 is crucial for our approach to solving the bidding problem (C), yet it is likely to sacrifice optimality. Indeed, selecting the water levels \mathbf{w}_d already at the end of day $d-1$ implies that \mathbf{w}_d cannot adapt to the sequence of reserve allocations $\{(\kappa_t^u, \kappa_t^v)\}_{t \in \mathcal{T}(d)}$ and reserve activations $\{(\rho_t^u, \rho_t^v)\}_{t \in \mathcal{T}(d)}$ encountered during day d . Thus, the water levels \mathbf{w}_d need to be sufficiently low so as to accommodate for any possible allocation and activation sequence throughout day d , and any water saved due to the absence of reserve-up or the presence of reserve-down activations will not be

accounted for on subsequent days. While this may discourage bidding on the reserve markets, our case study in Section 6 will show that significant fractions of the overall revenues are nevertheless generated from participation on the reserve markets.

We will now show that the trader's problem (CT) collapses to a tractable linear program if the reservoir targets imposed by the planner are chosen a day in advance. We thus subject the planner's problem (CP) to the information restriction of Approximation 1, which reduces the achievable expected revenues but makes the feasible set of the trader's problem (CT) independent of all exogenous uncertainties except for the reserve allocations $\{(\kappa_t^u, \kappa_t^v)\}_{t \in \mathcal{T}(d)}$ and the reserve activations $\{(\rho_t^u, \rho_t^v)\}_{t \in \mathcal{T}(d)}$. As the objective function of (CT) depends only on the bidding decisions $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(d)}$, we may restrict the operational decisions $\{(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ to depend only on the reserve activations without sacrificing optimality. Doing so results in the following variant of the trader's problem,

$$\begin{aligned}
& \sup && \sum_{t \in \mathcal{T}(d)} \tilde{\pi}_t^s s_t + \tilde{\pi}_t^u u_t + \tilde{\pi}_t^v v_t \\
& \text{s.t.} && s_t, u_t, v_t \in \mathbb{R}, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]}^\rho) && \forall t \in \mathcal{T}(d) \\
& && 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t && \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& && s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t && \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& && \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t && \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& && \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) && \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.},
\end{aligned} \tag{3}$$

where $\mathcal{F}_{[\downarrow(d), t]}^\rho = \sigma(\{(\kappa_\tau^u, \kappa_\tau^v)\}_{\tau \in \mathcal{T}(d)}, (\rho_\tau^u, \rho_\tau^v)\}_{\tau=\downarrow(d)}^t)$ denotes the σ -algebra generated by all reserve allocations of day d as well as all reserve activations within the interval $[\downarrow(d), t]$. Even though (3) constitutes a restriction of the trader's problem (CT), their optimal values can be shown to coincide.

PROPOSITION 3. *Under Approximation 1, the optimal values of (CT) and (3) are equal.*

Approximation 1 thus reduces the trader's problem (CT) to the stochastic program (3), which accommodates only $H + 1$ decision stages and whose wait-and-see decisions depend only on the reserve allocations and activations. However, (3) still constitutes an infinite-dimensional linear program. We will now show that problem (3) is indeed equivalent to an efficiently solvable linear program. As a first step towards this goal, we derive a family of valid inequalities that may be appended to problem (3) without affecting its feasible set.

PROPOSITION 4. *Any feasible solution of (3) satisfies $s_t - \bar{\rho}_t^v v_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t$ for all $t \in \mathcal{T}(d)$.*

The valid inequalities of Proposition 4 characterize the maximum bids on the reserve-down market (collectively across all arcs of the reservoir system) that can be honored under all possible realizations of the reserve activations. To see this, note that $\zeta_{t,a}\bar{p}_{t,a}$ represents the maximum amount of energy that can be absorbed on arc a by pumping. In case of a call-off on the reserve-down market, the energy production on all arcs, combined, can therefore be reduced at most by $s_t + \sum_{a \in \mathcal{A}} \zeta_{t,a}\bar{p}_{t,a}$. Appending these valid inequalities to problem (3) yields

$$\begin{aligned}
& \sup \quad \sum_{t \in \mathcal{T}(d)} \tilde{\pi}_t^s s_t + \tilde{\pi}_t^u u_t + \tilde{\pi}_t^v v_t \\
& \text{s.t.} \quad s_t, u_t, v_t \in \mathbb{R}, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d),t]}^\rho) & \forall t \in \mathcal{T}(d) \\
& \quad 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t & \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\
& \quad s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t & \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\
& \quad s_t - \bar{\rho}_t^v v_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t & \forall t \in \mathcal{T}(d) \\
& \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t & \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\
& \quad \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) & \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.},
\end{aligned} \tag{4}$$

which is equivalent to problem (3) by Proposition 4.

Next, we show that problem (4) is equivalent to the reduced stochastic program

$$\begin{aligned}
& \sup \quad \sum_{t \in \mathcal{T}(d)} \tilde{\pi}_t^s s_t + \tilde{\pi}_t^u u_t + \tilde{\pi}_t^v v_t \\
& \text{s.t.} \quad s_t, u_t, v_t \in \mathbb{R}, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathbb{R}^A & \forall t \in \mathcal{T}(d) \\
& \quad 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t & \forall t \in \mathcal{T}(d) \\
& \quad s_t + \bar{\rho}_t^u u_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t & \forall t \in \mathcal{T}(d) \\
& \quad s_t - \bar{\rho}_t^v v_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t & \forall t \in \mathcal{T}(d) \\
& \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t & \forall t \in \mathcal{T}(d) \\
& \quad \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau).
\end{aligned} \tag{CT^r}$$

Problem (CT^r) involves only here-and-now decisions, and its constraints are deterministic conditional on $\boldsymbol{\xi}_{[\downarrow(d)]}$. In fact, it constitutes a linear program of size $\mathcal{O}(A \cdot H)$.

PROPOSITION 5. *Under Approximation 1, the optimal value of problem (CT^r) is non-inferior to that of problem (4).*

Note that Proposition 5 does not prove the equivalence of the problems (CT^r) and (4). Indeed, proving this equivalence directly appears to be difficult. Instead, Appendix C proves that under

Approximation 1, problem (CT) is a restriction of problem (CT^r). Since Propositions 3–5 also imply that (CT^r) is a restriction of (CT), both problems indeed share the same optimal value. The relations between the different variants of the trader’s problem are illustrated in Figure 1. Every arc encodes a relation $A \leq B$, where A and B represent the optimal values of the problems at the arc’s tail and head, respectively. Dashed arcs indicate trivial relaxations, and solid arcs represent non-trivial implications proved in the referenced propositions. Interpreting Figure 1 as a directed graph, we note that Proposition 5 and Appendix C complete a counter-clockwise loop that visits each node, which implies that the optimal values of *all* optimization problems are in fact equal.

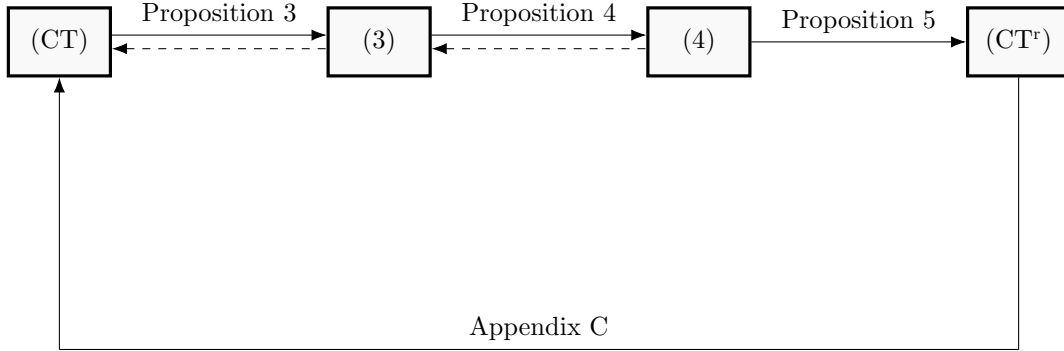


Figure 1 Illustration of the relations between different variants of the trader’s problem. Dashed arcs represent trivial relaxations, and solid arcs represent non-trivial implications proved in the referenced propositions.

The results of Section 4 as visualized in Figure 1 culminate in the following main theorem.

THEOREM 1. *Under Approximation 1, the optimal values of (CT) and (CT^r) are equal.*

In the remainder of this section we leverage Theorem 1 to construct a restriction of the bidding model (C) that is susceptible to further approximations and numerical solution. To this end, note first that Approximation 1 restricts the planner’s problem (CP) to

$$\begin{aligned}
 & \sup \sum_{d \in \mathcal{D}} \mathbb{E}[\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})] \\
 & \text{s.t. } \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{[\downarrow(d)]}) \quad \forall d \in \mathcal{D} \\
 & \quad \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.}
 \end{aligned} \tag{CP^r}$$

where the end-of-day reservoir levels \mathbf{w}_d are chosen one day in advance and thus adapt to information available at time $\downarrow(d)$. Under this information restriction, Theorem 1 allows us to compute

the optimal value $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\Psi(d)]})$ of the trader's problem (CT) by solving the linear program (CT^r), which enables us to prove that the planner's reduced problem (CP^r) is equivalent to

$$\begin{aligned}
& \sup \sum_{t \in \mathcal{T}} \mathbb{E} [\tilde{\pi}_t^s s_t + \tilde{\pi}_t^u u_t + \tilde{\pi}_t^v v_t] \\
\text{s.t. } & s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[T]}), \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[T]}), \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{[\Psi(d)]}) \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d) \\
& 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\
& s_t + \bar{\rho}_t^u u_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t, s_t - \bar{\rho}_t^v v_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \quad (\text{C}^r) \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall d \in \mathcal{D}, \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\
& \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.} \\
& \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.}
\end{aligned}$$

Problem (C^r) can be viewed as a reduction of the original bidding model (C).

THEOREM 2. *The optimal values of the problems (CP^r) and (C^r) are equal, and they are smaller than or equal to the optimal value of problem (C).*

The bidding model (C) and its reduction (C^r) differ in the following aspects.

- In problem (C) the market decisions (s_t, u_t, v_t) are adapted to information that is available at the *beginning of day* $d(t)$, while the operational decisions $(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)$ are adapted to *real-time information*. In contrast, in (C^r) these decisions are taken under *perfect information*.
- Problem (C) has *random recourse* because the reserve market bids in the energy delivery constraints are multiplied by the uncertain reserve activations. In contrast, problem (C^r) has *fixed recourse* because all reserve activations were eliminated from the constraints.
- Problem (C^r) accommodates the *valid cuts* derived in Proposition 4, which are absent in (C).

We highlight again that problem (C^r) was obtained from the original bidding model (C) by applying a single information restriction, whereby the end-of-day reservoir levels must be chosen one day in advance. No other approximations have been applied. Finally, the equivalence between (CP^r) and (C^r) established in Theorem 2 is in terms of the problems' optimal values and not necessarily of their optimal solutions. In essence, the proof shows that the optimal market decisions s_t, u_t and v_t as well as the optimal operational decisions $\mathbf{g}_t, \mathbf{p}_t$ and \mathbf{z}_t of problem (C^r), which may fail to be non-anticipative, can be converted to non-anticipative decisions

$$\begin{aligned}
s'_t &= \mathbb{E} [s_t | \boldsymbol{\xi}_{[\Psi(t)]}], & u'_t &= \mathbb{E} [u_t | \boldsymbol{\xi}_{[\Psi(t)]}], & v'_t &= \mathbb{E} [v_t | \boldsymbol{\xi}_{[\Psi(t)]}] \\
\mathbf{g}'_t &= \mathbb{E} [\mathbf{g}_t | \boldsymbol{\xi}_{[\Psi(t)]}], & \mathbf{p}'_t &= \mathbb{E} [\mathbf{p}_t | \boldsymbol{\xi}_{[\Psi(t)]}], & \mathbf{z}'_t &= \mathbb{E} [\mathbf{z}_t | \boldsymbol{\xi}_{[\Psi(t)]}]
\end{aligned}$$

by taking conditional expectations. Because the objective and constraint functions are all linear in these decision variables, the new solution remains feasible and—in fact—optimal in problem (C^r). This argument is only valid because Approximation 1 ensures that problem (C^r) has fixed recourse.

5. Numerical Solution of the Reduced Bidding Model

Although Approximation 1 is at the core of our reduction of the bidding model (C) to (C^r), it is not entirely innocent as it may discourage bidding on the reserve markets (*cf.* Section 4). Moreover, the reduced planner’s problem (CP^r) constitutes a large-scale stochastic program that is challenging to solve. To simultaneously alleviate both shortcomings, this section develops variants of the reduced planner’s and trader’s problems (CP^r) and (CT^r) with the following characteristics: (i) the reduced planner’s problem employs an affine decision rule approximation to result in a two-stage stochastic program that can be solved efficiently using the well-known sample average approximation; (ii) the reduced trader’s problem accounts not only for the worst possible reserve activation sequence but also for the nominal activation sequence; and (iii) the reduced trader’s problem contains a valuation of the nominal end-of-day reservoir levels in the objective function. We also incorporate both problems into a rolling horizon framework that repeatedly solves the planner’s and trader’s problems and always only implements the decisions of the first day. Under these changes, the equivalences of Theorems 1 and 2 no longer hold, but we expect the resulting variants of problems (CP^r) and (CT^r) to perform better in actual case studies (*cf.* Section 6).

In the following, we first describe how to efficiently solve and extract water values from problem (C^r). Afterwards, we discuss how these water values can be used to compute the first-day decisions by solving a variant of the trader’s problem that accounts both for the water target levels and for valuations of the end-of-day water levels.

In order to compute water values, we consider an equivalent reformulation of the reduced planner’s problem (CP^r) that contains the superfluous decision variables \mathbf{w}_d^- , $d \in \mathcal{D}$, to denote the water levels at the *beginning* of day $d \in \mathcal{D}$:

$$\begin{aligned} & \sup \sum_{d \in \mathcal{D}} \mathbb{E}[\Pi_d(\mathbf{w}_d^-, \mathbf{w}_d, \boldsymbol{\xi}_{\lfloor \Psi(d) \rfloor})] \\ \text{s.t. } & \mathbf{w}_d^- \in \mathcal{L}^R(\mathcal{F}_{\lfloor \Psi(d-1) \rfloor}), \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{\lfloor \Psi(d) \rfloor}) \quad \forall d \in \mathcal{D} \\ & \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.} \\ & \mathbf{w}_d^- = \mathbf{w}_{d-1} \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.} \end{aligned} \quad (5)$$

The shadow prices of the newly introduced constraints $\mathbf{w}_d^- = \mathbf{w}_{d-1}$ can be interpreted as water values at the end of day $d-1$, and they will be used later on in our variant of the trader’s problem.

While the reduced bidding model (5) has fixed recourse and takes all market bids and operational decisions under perfect information, it still constitutes an infinite-dimensional linear program over a space of measurable functions, and thus there is little hope to solve it *exactly*. To make this problem amenable to numerical solution, we further reduce it to a two-stage stochastic program by restricting the reservoir filling levels to parsimonious *affine decision rules* that depend on the observable random parameters only through a few judiciously chosen features.

Approximation 2 (Affine Decision Rule Restriction) *For all $d \in \mathcal{D}$, the end-of-day reservoir filling levels are representable as*

$$\mathbf{w}_d = \boldsymbol{\lambda}_d + \boldsymbol{\Phi}_d \left(\sum_{\tau \in \mathcal{T}(d)} \phi_\tau \right) + \boldsymbol{\Psi}_d \left(\sum_{\tau=1}^{\uparrow(d-1)} \phi_\tau \right) + \boldsymbol{\mu}_d \left(\frac{1}{H} \sum_{\tau \in \mathcal{T}(d)} \pi_\tau^s \right)$$

for some fixed vectors $\boldsymbol{\lambda}_d, \boldsymbol{\mu}_d \in \mathbb{R}^R$ and matrices $\boldsymbol{\Phi}_d, \boldsymbol{\Psi}_d \in \mathbb{R}^{R \times R}$.

In principle, the decision rule for \mathbf{w}_d can depend on any linear transformation of $\boldsymbol{\xi}_{[\downarrow(d)]}$. Generally, richer feature vectors lead to more adaptive (and hence better) decisions at a higher computational cost. We do not include as features the reserve allocations and activations since they are assumed to be serially independent as well as independent of all other exogenous uncertainties. Instead, Approximation 2 restricts the reservoir filling levels at the end of day d to affine functions of the following features: (i) the cumulative natural inflows into the reservoirs across day d , (ii) the cumulative natural inflows into the reservoirs across the planning horizon until the end of day $d-1$ and (iii) the average spot price on day d . The proposed affine decision rules are parsimonious as they compress the history of all observations into a few relevant features, but they are flexible enough to allow the planner to set different reservoir targets for wet and dry days, for wet and dry seasons as well as for high- and low-price days. All of the proposed features are observable at the beginning of day d , and therefore the non-anticipativity conditions $\mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{[\downarrow(d)]})$, $d \in \mathcal{D}$, imposed by Approximation 1 are automatically satisfied. Affine decision rule approximations are routinely used in multistage robust optimization and stochastic programming problems; we refer to the reviews by Ben-Tal et al. (2009), Delage and Iancu (2015) and Yanıkoğlu et al. (2019).

Note that Approximation 2 restricts the functional form of the reservoir targets, and thus it results in a conservative lower bound on the optimal value of problem (C^r), which itself underestimates the optimal value of the original bidding model (C) by virtue of Theorem 2. We further emphasize that Approximation 2 reduces problem (5) to a two-stage stochastic program with here-and-now decisions $\boldsymbol{\lambda}_d, \boldsymbol{\mu}_d \in \mathbb{R}^R$ and $\boldsymbol{\Phi}_d, \boldsymbol{\Psi}_d \in \mathbb{R}^{R \times R}$, $d \in \mathcal{D}$, which are chosen without any information about $\boldsymbol{\xi}_{[T]}$, and with wait-and-see decisions $s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[T]})$ and $\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[T]})$, $t \in \mathcal{T}$, which are chosen under perfect information about $\boldsymbol{\xi}_{[T]}$. The emerging two-stage stochastic program can then be solved with the popular sample average approximation (Shapiro et al. 2009).

Approximation 3 (Sample Average Approximation) *The original probability measure \mathbb{P} is replaced with a discrete empirical measure $\widehat{\mathbb{P}} = \frac{1}{N} \sum_{n=1}^N \delta_{\omega^{(n)}}$, where $\delta_{\omega^{(n)}}$ stands for the Dirac point mass at $\omega^{(n)}$, and where $\omega^{(n)} \in \Omega$, $n = 1, \dots, N$, constitute independent samples from \mathbb{P} .*

The combination of affine decision rules and the sample average approximation is not new and has been explored, among others, by Vayanos et al. (2012) and Bodur and Luedtke (2022) in the

context of robust optimization and stochastic programming, respectively. We emphasize that, even though we use an empirical probability measure $\widehat{\mathbb{P}}$ when solving the decision rule approximation of problem (C^r), the conditional expectations $\tilde{\pi}_t^s$, $\tilde{\pi}_t^u$ and $\tilde{\pi}_t^v$ in the objective function of (C^r) are pre-computed under the original measure \mathbb{P} . We solve the linear program obtained by applying Approximations 2 and 3 to problem (5) and record, in addition to the water target levels \mathbf{w}_1 , the water values $\boldsymbol{\vartheta}_1 \in \mathbb{R}^R$ corresponding to the water preservation constraints $\mathbf{w}_2^- = \mathbf{w}_1$. Note that these water values are deterministic since \mathbf{w}_1 and \mathbf{w}_2^- are deterministic by virtue of Approximation 1, and that $\vartheta_{1,R} = 0$ by the properties of the dummy reservoir R .

To compute the market bids $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(1)}$ for day $d = 1$, we solve the following variant of the trader's problem,

$$\begin{aligned}
& \sup \sum_{t \in \mathcal{T}(1)} \tilde{\pi}_t^s s_t + \tilde{\pi}_t^u u_t + \tilde{\pi}_t^v v_t + \boldsymbol{\vartheta}_1^\top [\boldsymbol{\phi}_t + \mathbf{M}(\hat{\mathbf{g}}_t - \hat{\mathbf{p}}_t + \hat{\mathbf{z}}_t)] \\
& \text{s.t. } s_t, u_t, v_t \in \mathbb{R}, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathbb{R}^A, \hat{\mathbf{g}}_t, \hat{\mathbf{p}}_t, \hat{\mathbf{z}}_t \in \mathbb{R}^A \quad \forall t \in \mathcal{T}(1) \\
& 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t, \hat{\mathbf{g}}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t, \hat{\mathbf{p}}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t, \hat{\mathbf{z}}_t \quad \forall t \in \mathcal{T}(1) \\
& s_t + \bar{\rho}_t^u u_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t, s_t - \bar{\rho}_t^v v_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t \quad \forall t \in \mathcal{T}(1) \\
& s_t + \hat{\kappa}_t^u \hat{\rho}_t^u u_t - \hat{\kappa}_t^v \hat{\rho}_t^v v_t = \boldsymbol{\eta}_t^\top \hat{\mathbf{g}}_t - \boldsymbol{\zeta}_t^\top \hat{\mathbf{p}}_t \quad \forall t \in \mathcal{T}(1) \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_0 + \sum_{\tau=\downarrow(1)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(1) \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_0 + \sum_{\tau=\downarrow(1)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\hat{\mathbf{g}}_\tau - \hat{\mathbf{p}}_\tau + \hat{\mathbf{z}}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(1) \\
& \mathbf{w}_1 \leq \mathbf{w}_0 + \sum_{\tau \in \mathcal{T}(1)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau),
\end{aligned} \tag{6}$$

where $\hat{\kappa}_t^u = \mathbb{E}[\kappa_t^u]$, $\hat{\kappa}_t^v = \mathbb{E}[\kappa_t^v]$, $\hat{\rho}_t^u = \mathbb{E}[\rho_t^u]$ and $\hat{\rho}_t^v = \mathbb{E}[\rho_t^v]$ for all $t \in \mathcal{T}(1)$. Problem (6) differs from the day-1 reduced trader's problem (CT^r) of Section 4 in two aspects. Firstly, the operational decisions $\{(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(1)}$ guaranteeing the satisfaction of the reservoir bounds under any possible reserve activation sequence are complemented with the *nominal* operational decisions $\{(\hat{\mathbf{g}}_t, \hat{\mathbf{p}}_t, \hat{\mathbf{z}}_t)\}_{t \in \mathcal{T}(1)}$ that account for the water dynamics under the expected scenario, where the products of the reserve allocations and activations are replaced with their means $\mathbb{E}[\kappa_t^u \rho_t^u] = \hat{\kappa}_t^u \hat{\rho}_t^u$ and $\mathbb{E}[\kappa_t^v \rho_t^v] = \hat{\kappa}_t^v \hat{\rho}_t^v$ for all $t \in \mathcal{T}(1)$. Secondly, problem (6) incorporates both the water target level constraints and a valuation of the nominal end-of-day reservoir levels in the objective function. These changes alleviate the conservatism introduced by Approximation 1, but they imply that the equivalence of Theorem 1 no longer holds for the revised trader's problem (6).

To obtain the operational decisions $(\mathbf{g}_\theta, \mathbf{p}_\theta, \mathbf{z}_\theta)$ for each hour $\theta \in \mathcal{T}(1)$ that are compatible with the market bids (s_t, u_t, v_t) and feasible in view of the revealed reserve market activations, finally, we solve the truncated problems

$$\begin{aligned}
& \sup \quad \boldsymbol{\vartheta}_1^\top \mathbf{M} \left(\mathbf{g}_\theta - \mathbf{p}_\theta + \mathbf{z}_\theta + \sum_{t=\theta+1}^H (\hat{\mathbf{g}}_t - \hat{\mathbf{p}}_t + \hat{\mathbf{z}}_t) \right) \\
& \text{s.t.} \quad \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathbb{R}^A, \hat{\mathbf{g}}_t, \hat{\mathbf{p}}_t, \hat{\mathbf{z}}_t \in \mathbb{R}^A \quad \forall t \in [\theta, H] \\
& \quad \mathbf{0} \leq \mathbf{g}_t, \hat{\mathbf{g}}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t, \hat{\mathbf{p}}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t, \hat{\mathbf{z}}_t \quad \forall t \in [\theta, H] \\
& \quad s_\theta + \kappa_\theta^u \rho_\theta^u u_\theta - \kappa_\theta^v \rho_\theta^v v_\theta = \boldsymbol{\eta}_\theta^\top \mathbf{g}_\theta - \boldsymbol{\zeta}_\theta^\top \mathbf{p}_\theta \\
& \quad s_t + \bar{\rho}_t^u u_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t, \quad \forall t \in [\theta + 1, H] \\
& \quad s_t + \kappa_t^u \hat{\rho}_t^u u_t - \kappa_t^v \hat{\rho}_t^v v_t = \boldsymbol{\eta}_t^\top \hat{\mathbf{g}}_t - \boldsymbol{\zeta}_t^\top \hat{\mathbf{p}}_t \quad \forall t \in [\theta + 1, H] \\
& \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_0 + \sum_{\tau=1}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in [\theta, H] \\
& \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_0 + \sum_{\tau=1}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\hat{\mathbf{g}}_\tau - \hat{\mathbf{p}}_\tau + \hat{\mathbf{z}}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in [\theta, H] \\
& \quad \mathbf{w}_1 \leq \mathbf{w}_0 + \sum_{\tau \in \mathcal{T}(1)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau).
\end{aligned}$$

Each of these problems differs from problem (6) in two aspects. Firstly, the truncated problems disregard the impact of the (sunk) market bids $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(1)}$ on the objective function. Secondly, the nominal reserve allocations and activations $\hat{\kappa}_t^u, \hat{\kappa}_t^v, \hat{\rho}_\theta^u$ and $\hat{\rho}_\theta^v$ in hour θ are replaced with their actual realizations $\kappa_t^u, \kappa_t^v, \rho_\theta^u$ and ρ_θ^v , respectively, for all $t \in [\theta + 1, H]$.

Apart from reducing the conservatism introduced by Approximation 1, our rolling horizon implementation enjoys several advantages: (i) The true state of the world ω may differ from all discretization points $\omega^{(n)}$, $n = 1, \dots, N$, of the empirical measure $\hat{\mathbb{P}}$, and thus the sample average approximation does not provide any recourse decisions corresponding to ω . (ii) Problem (C) only models a finite time window of the perpetual operation of the reservoir system, which adversely affects the decisions towards the end of the planning horizon \mathcal{D} . (iii) Resolving problem (C) on a daily basis ensures that the end-of-day reservoir targets adapt to all available information, and it reduces the conservatism of the affine decision rule approximation.

6. Case Study: Gasteiner Tal Cascade

We apply our planner-trader decomposition to a hydropower cascade located in the Gasteiner Tal, Austria. We describe the problem instance in Section 6.1, and we present our results in Section 6.2.

6.1. Problem Parameters

The Gasteiner Tal cascade comprises three reservoirs: the Bockhartsee annual reservoir with a capacity of 18,500,000 m³, the Nassfeld daily reservoir with a capacity of 230,000 m³ and the Remsach compensation reservoir with a capacity of 4,000 m³. The Bockhartsee reservoir is connected to the Nassfeld reservoir by the Nassfeld pumped-storage plant that has a generating capacity of 40,600 m³/h and a generation efficiency of 6.68×10^{-4} MWh/m³, as well as a pumping capacity of 28,500 m³/h and an inverse pumping efficiency of 9.35×10^{-4} MWh/m³, respectively. The Nassfeld reservoir is connected to the Remsach reservoir by the Böckstein plant with a generating capacity of 41,400 m³/h and an efficiency of 10.3×10^{-4} MWh/m³, and the Remsach reservoir is connected to the Gasteiner Ache river by the Remsach plant with a generating capacity of 50,400 m³/h and an efficiency of 5.41×10^{-4} MWh/m³. In total, the cascade produces up to 264,000 MWh of electricity per year and covers the demand of approximately 75,000 households.¹

We model the hourly natural inflows into the three reservoirs based on historical daily inflow data for the years 2002-2012. In absence of more detailed data, we make the simplifying assumption that each hourly inflow within a day amounts to 1/24-th of the corresponding daily inflow. We use the inflows of the years 2002-2011 as in-sample scenarios and the inflows of the year 2012 as out-of-sample scenarios, respectively.

For the spot and reserve prices, we use the Balancing Statistics dataset for the year 2020 that is publicly available on the website of the Austrian Power Grid AG.² We use the hourly averages of the EXAA and EPEX spot prices as the out-of-sample spot prices, whereas our in-sample spot prices are generated from a linear regression model of the out-of-sample prices with indicator variables for the month of the year, the day of the week and the hour of the day. In contrast, the capacity fees and the activation prices on the secondary reserve markets are reported in four-hourly and 15-minute blocks, respectively. We use the four-hourly capacity fees for the secondary reserve-up and reserve-down market as the out-of-sample capacity fees, while our in-sample capacity fees are derived from a logarithmic regression model of the corresponding realized fees with indicator variables for the month of the year, the day of the week and the four-hour block of the day. Likewise, we use the realized activation prices on the reserve markets at 15-minute granularity as out-of-sample data, whereas the in-sample prices are generated from a linear regression model of the out-of-sample prices with indicator variables for the month of the year, the day of the week and the hour of the day. While our regression models appear to generally perform reasonably well, we note that the capacity fees appear to be challenging to predict to a high accuracy. We also note that despite the

¹ Further details of the cascade can be found online at <https://www.salzburg-ag.at/content/dam/web18/dokumente/unternehmen/erzeugung/Kraftwerke-GasteinerTal.pdf>.

² Website: www.apg.at.

	0.0	0.25	0.5	0.75	1.0
capacity	(28.7%, 26.9%)	(39.4%, 32.8%)	(42.4%, 38.7%)	(67.6%, 46.7%)	(70.7%, 47.1%)
energy	(0.9%, 1.6%)	(24.4%, 26.8%)	(50.9%, 51.8%)	(78.0%, 77.9%)	(99.3%, 99.1%)

Table 1 Rates at which capacity bids (first row) and energy price bids (second row) are accepted for different target probabilities (columns) on the reserve-up (first entry in each cell) and the reserve-down (second entry in each cell) markets.

availability of further historical prices, we decided to use the year 2020 as the basis for both the in-sample and the out-of-sample data as the energy prices exhibited significant non-stationarity in recent years due to the changing political climate. Note, however, that the in-sample prices differ from the out-of-sample prices due to our use of the aforementioned regression models.

6.2. Results

We now apply our decomposition scheme to the case study of Section 6.1. Our experiments comprise a planning horizon of one year (52 weeks \times 7 days, resulting in $D = 364$ days). To alleviate end-of-horizon effects, we solve our master problem as a rolling-horizon problem where the in-sample and out-of-sample data is re-used in a circular fashion, that is, we re-use the data of year 2020 *in lieu* of the data for the subsequent year 2021. To account for the idiosyncracies of the energy market of our case study, we also incorporate constraints into our planner and trader problems that require the hourly capacity fee and activation price bids to be constant across 4-hour blocks. Our planner and trader problems are controlled by two hyper-parameters that specify the target probability with which our capacity bids are being accepted as well as the target frequency with which our energy price bids are being activated, respectively. The actual capacity bids and energy prices vary across the four-hour blocks, and they are constructed from the hyper-parameters using our regression models for the capacity fees and activation prices, respectively. Below, we present results for the hyper-parameter combinations $\{0.0, 0.25, 0.5, 0.75, 1.0\} \times \{0.0, 0.25, 0.5, 0.75, 1.0\}$. Our decomposition scheme was implemented in C++ using Gurobi 9.0.1 and run on an 3.2 GHz Intel Xeon processor in single-core mode. All planner problems were solved within less than two minutes, and the trader problems were all solved within fractions of a second. For the sake of exposition, all water levels and revenues are denominated in 1,000s of m³ and 1,000s of EUR, respectively. Our datasets and detailed results, together with the source codes of all algorithms, can be found online.³

We first consider the choice of the two hyper-parameters. To this end, Table 1 compares the target acceptance probabilities for our capacity bids with the actual acceptance frequencies (first row) as well as the target activation frequencies for our energy price bids with the actual activation frequencies (second row) on the reserve-up (first entry) and reserve-down (second entry) markets.

³ Website: <https://github.com/napat-rujeerapaiboon/hydro-scheduling>.

	0.0	0.25	0.5	0.75	1.0
0.0	10,755.29 (100.0%, 0.0%, 0.0%)	10,761.46 (99.9%, 0.1%, 0.0%)	10,769.91 (99.7%, 0.2%, 0.0%)	10,771.26 (99.8%, 0.2%, 0.0%)	10,767.04 (99.8%, 0.2%, 0.0%)
0.25	11,660.82 (99.3%, 0.0%, 0.6%)	11,065.15 (94.6%, 1.5%, 3.9%)	11,507.15 (94.0%, 3.1%, 2.9%)	11,256.31 (95.2%, 4.0%, 0.8%)	10,842.81 (103.3%, 2.0%, -5.2%)
0.5	10,479.21 (97.7%, 0.4%, 1.8%)	10,722.82 (95.2%, 0.0%, 4.8%)	11,800.98 (87.4%, 9.4%, 3.2%)	11,424.93 (89.8%, 10.6%, -0.4%)	10,539.76 (95.5%, 10.7%, -6.2%)
0.75	10,932.94 (96.9%, 0.4%, 2.7%)	12,237.71 (83.1%, 11.9%, 5.0%)	12,836.11 (74.5%, 22.6%, 2.9%)	12,700.91 (63.9%, 38.0%, -1.9%)	11,077.83 (67.2%, 41.9%, -9.1%)
1.0	12,717.43 (97.6%, 0.5%, 2.0%)	12,130.75 (79.2%, 16.0%, 4.7%)	13,396.20 (53.0%, 44.8%, 2.2%)	13,268.93 (93.9%, 6.5%, -0.4%)	13,141.65 (67.2%, 40.8%, -8.0%)

Table 2 Revenues (top entry) and revenue contributions by market (bottom entry; from left to right: spot, reserve-up and reserve-down market) for different combinations of target probabilities for the capacity bids (rows) and energy price bids (columns).

The table shows that while our predictions of the energy price activation frequencies are remarkably accurate, the actual capacity bid acceptance frequencies can deviate quite significantly from the target acceptance probabilities. We attribute this to the relatively poor fit of the underlying vanilla regression models described in the previous section. Table 2 compares the out-of-sample revenues generated by different combinations of both hyper-parameters over the course of the year, and it also elucidates how the spot and the two reserve markets contribute towards those revenues. The table reveals that overall, targeting a 100% probability of having the capacity bids accepted (fifth row) and targeting for a 50% frequency of having the energy price bids accepted (third row), respectively, performs best. Note that while almost half of the revenues generated under this hyper-parameter combination stem from the reserve-up market, the reserve-down market appears to contribute only a negligible fraction to the overall revenues. This turns out to be misleading, however, as successful bids on the reserve-down market lead to water being pumped to higher reservoirs, which implies further revenue uplifts at later times that we do not account for in this calculation. This is further evidenced by separate numerical experiments (not reported here) which reveal that the inclusion of either reserve market to a spot-only model leads to a revenue uplift of 20%-25%, whereas the additional inclusion of the other reserve market only provides a further uplift of 0.1%-5%, independent of the order in which the two reserve markets are added.

We now take a closer look at the best-performing strategy from Table 2. Figure 2 visualizes the cumulative contributions of the spot and reserve markets towards the energy bids (left), the energy provided (middle) and the revenues generated (right). The figure shows that while the spot and the reserve-up market both generate about half of the overall revenues, the amounts of energy exchanged on all three markets are of the same order of magnitude. Figures 3 and 4 present the water levels and water values for the three reservoirs of our case study, respectively. The figures show that our model correctly captures the different time scales of the reservoir dynamics. While the filling level of the Bockhartsee annual reservoir changes slowly throughout the year, the Nassfeld

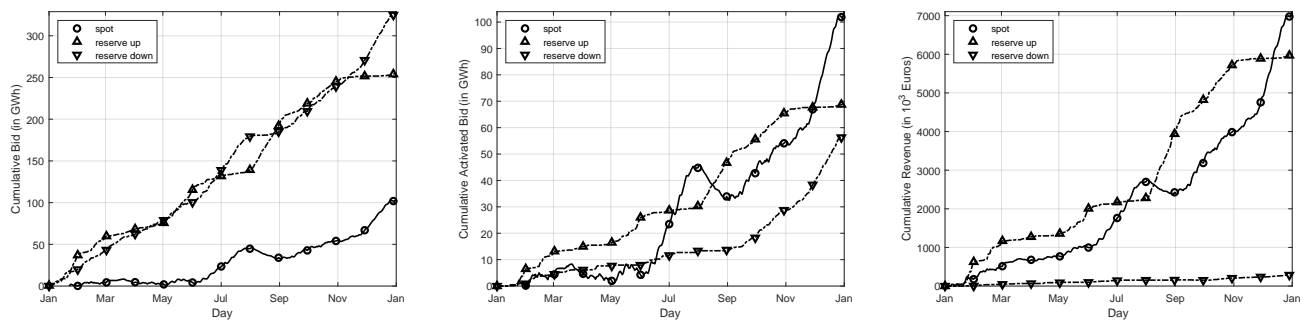


Figure 2 Cumulative bids and revenue contributions of the spot and the reserve markets. From left to right, the graphs show the cumulative bids, the activated bids and the generated revenues on each of the three markets.




Figure 3 Out-of-sample reservoir levels for the Bockhartsee (left), Nassfeld (middle) and Remsach (right) reservoirs.

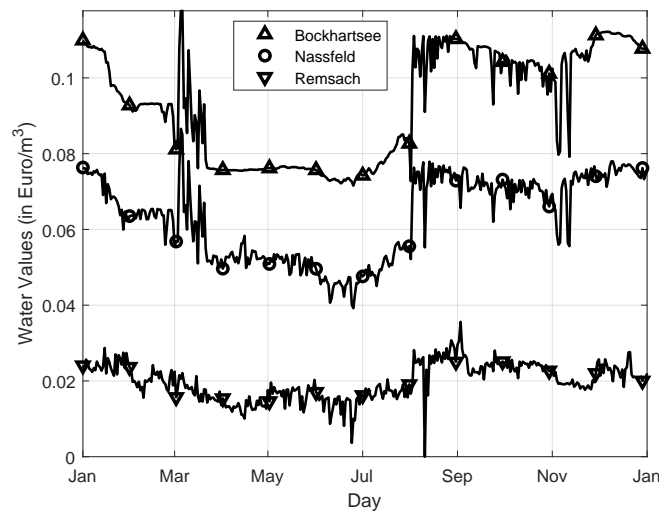


Figure 4 Water values for the Bockhartsee (left), Nassfeld (middle) and Remsach (right) reservoirs.

and Remsach reservoirs operate at a much faster pace. The seasonality of the water levels also implies a seasonality in the water values, which are highest ahead of the depletion phase of the top-level reservoir. Note that in our experiments, all lower reservoir bounds are set to zero. In

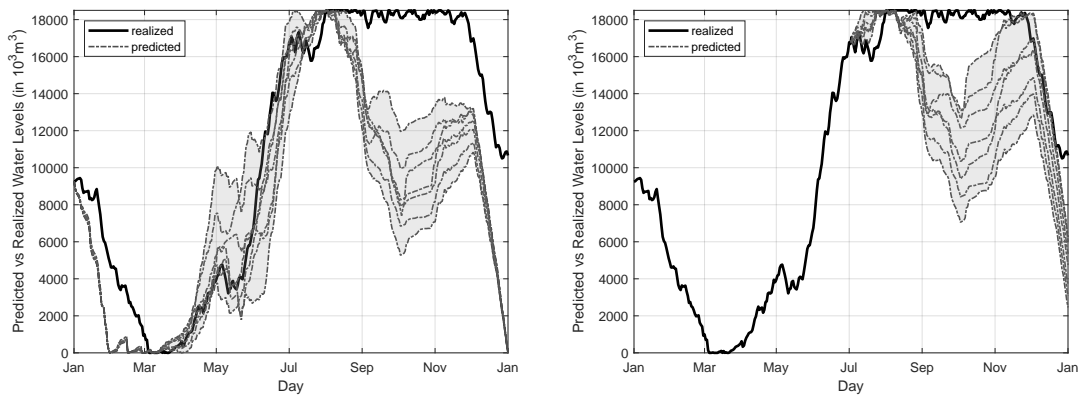


Figure 5 Predictive accuracy of the planner's problem. The realized out-of-sample water levels of the Bockhartsee reservoir are printed in bold, whereas the shaded regions illustrate the in-sample water level predictions on day 1 (left graph) and on day 182 (right graph).

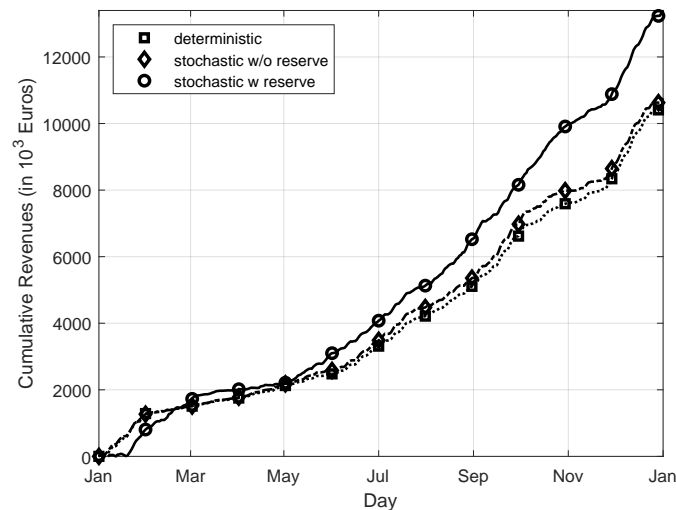


Figure 6 Cumulative revenues generated by the deterministic model, the stochastic model without reserve markets and the stochastic model including reserve markets.

practice, one typically imposes strictly positive lower reservoir bounds to protect the aquatic life as well as to avoid a negative impact on tourism. Figure 5 investigates the predictive accuracy of our bidding model. It visualizes the out-of-sample water levels of the Bockhartsee reservoir across the time horizon, together with the in-sample prediction of these water levels on day 1 and day 182. We observe that the out-of-sample water levels tend to stay close the predicted intervals for a while, but that our predictions become increasingly inaccurate as we look further into the future. We emphasize that the water level predictions are updated every day, rather than twice during the time horizon as shown in this figure.

In our final experiment, we compare our bidding problem with a variant of the problem that precludes participation on the reserve markets. We also compare our model against a deterministic bidding formulation that replaces all stochastic parameters with their expected values. In particular, this formulation replaces the stochastic reserve market activations with deterministic fractional activations, which implies that the hydropower producer may not be able to honor her reserve market commitments out-of-sample. In such cases, we revert to a spot-only trading strategy for the affected days. We observe empirically that the cumulative out-of-sample revenues of the deterministic formulation are maximized when the total reserve market participation is restricted to not exceed the spot market participation in any given hour. This ad hoc restriction avoids overly aggressive reserve market participation that underperforms out-of-sample due to frequent switches to the aforementioned spot-only strategy. Note that by retrospectively switching to a spot-only strategy based on the observed out-of-sample reserve activations, the deterministic strategy has an advantage and is not implementable in reality. Figure 6 visualizes the cumulative out-of-sample revenues generated by the three models. The figure highlights the benefits of a stochastic strategy that participates in both the spot and the reserve markets: our strategy outperforms both the spot-only stochastic model and the deterministic model by about 25% in terms of cumulative revenues.

Acknowledgments

This publication is posthumously dedicated to the first author, Kilian Schindler (29.10.1989–16.5.2020), who has led the research in collaboration with the second author. We are grateful to Georg Ostermaier from Decision Trees GmbH and Dirk Lauinger for valuable discussions and advice, as well as the associate editor and the referees for their suggestions that significantly improved the manuscript. This research was supported by the Ministry of Education, Singapore, under its 2019 Academic Research Fund Tier 3 grant MOE-2019-T3-1-010 and the Swiss National Science Foundation grant BSCGI0_157733 as well as the EPSRC grant EP/R045518/1.

References

- Aasgård, E., G. Andersen, S. Fleten, D. Haugstvedt. 2014. Evaluating a stochastic-programming-based bidding model for a multireservoir system. *IEEE Transactions on Power Systems* **29**(4) 1748–1757.
- Aasgård, E., S. Fleten, M. Kaut, K. Midthun, G. Perez-Valdes. 2019. Hydropower bidding in a multi-market setting. *Energy Systems* **10** 543–565.
- Abgottspon, H., G. Andersson. 2012. Approach of integrating ancillary services into a medium-term hydro optimization. *Proceedings of the 12th Symposium of Specialists in Electric Operational and Expansion Planning*. 1–10.
- Abgottspon, H., K. Njålsson, M. A. Bucher, G. Andersson. 2014. Risk-averse medium-term hydro optimization considering provision of spinning reserves. *Proceedings of the 2014 International Conference on Probabilistic Methods Applied to Power Systems*. 1–6.

- Andrieu, L., R. Henrion, W. Römisch. 2010. A model for dynamic chance constraints in hydro power reservoir management. *European Journal of Operational Research* **207**(2) 579–589.
- Ban, G.-Y., C. Rudin. 2019. The big data newsvendor: Practical insights from machine learning. *Operations Research* **67**(1) 90–108.
- Ben-Tal, A., L. El Ghaoui, A. Nemirovski. 2009. *Robust Optimization*. Princeton University Press.
- Bertsimas, D., N. Kallus. 2020. From predictive to prescriptive analytics. *Management Science* **66**(3) 1025–1044.
- Bodur, M., J. R. Luedtke. 2022. Two-stage linear decision rules for multi-stage stochastic programming. *Mathematical Programming* **191**(1) 347–380.
- Carpentier, P.-L., M. Gendreau, F. Bastin. 2013. Long-term management of a hydroelectric multireservoir system under uncertainty using the progressive hedging algorithm. *Water Resources Research* **49**(5) 2812–2827.
- Chazarra, M., J. García-González, J. Pérez-Díaz, M. Arteseros. 2016. Stochastic optimization model for the weekly scheduling of a hydropower system in day-ahead and secondary regulation reserve markets. *Electric Power Systems Research* **130** 67–77.
- Chazarra, M., J. I. Pérez-Díaz, J. García-González. 2018. Optimal joint energy and secondary regulation reserve hourly scheduling of variable speed pumped storage hydropower plants. *IEEE Transactions on Power Systems* **33**(1) 103–115.
- Chazarra, M., J. I. Pérez-Díaz, J. García-González. 2014. Optimal operation of variable speed pumped storage hydropower plants participating in secondary regulation reserve markets. *Proceedings of the 11th International Conference on the European Energy Market (EEM14)*. 1–5.
- Chazarra, M., J. I. Pérez-Díaz, J. García-González, R. Praus. 2018. Economic viability of pumped-storage power plants participating in the secondary regulation service. *Applied Energy* **216** 224–233.
- de Matos, V., A. Philpott, E. Finardi. 2015. Improving the performance of stochastic dual dynamic programming. *Journal of Computational and Applied Mathematics* **290** 196–208.
- Delage, E., D. Iancu. 2015. Robust multistage decision making. *INFORMS Tutorials in Operations Research* 19–46.
- DOE. 2020. DOE OE global energy storage database. <https://www.sandia.gov/ess-ssl/global-energy-storage-database/>. Accessed on 1st March 2021.
- Dyer, M., L. Stougie. 2006. Computational complexity of stochastic programming problems. *Mathematical Programming* **106**(3) 423–432.
- European Commission. 2020. European Commission METIS technical note T4: Overview of European electricity markets. https://ec.europa.eu/energy/sites/default/files/documents/overview_of_european_electricity_markets.pdf. Accessed on 1st March 2021.

- Fernández-Muñoz, D., J. I. Pérez-Díaz, I. Guisández, M. Chazarra, Á. Fernández-Espina. 2020. Fast frequency control ancillary services: An international review. *Renewable and Sustainable Energy Reviews* **120** 109662.
- Fodstad, M., M. Aarlott, K. Midthun. 2018. Value-creation potential from multi-market trading for a hydropower producer. *Energies* **11**(1) 1–15.
- Gauvin, C., E. Delage, M. Gendreau. 2017. Decision rule approximations for the risk averse reservoir management problem. *European Journal of Operational Research* **261**(1) 317–336.
- Gauvin, C., E. Delage, M. Gendreau. 2018. A stochastic program with time series and affine decision rules for the reservoir management problem. *European Journal of Operational Research* **267**(2) 716–732.
- Gjelsvik, A., B. Mo, A. Haugstad. 2010. Long- and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming. P., Pardalos, S. Rebennack, M. Pereira, N. Iliadis, eds., *Handbook of Power Systems I*. Springer, 33–55.
- Hanasusanto, G., D. Kuhn, W. Wiesemann. 2016. A comment on “Computational complexity of stochastic programming problems”. *Mathematical Programming* **159** 557–569.
- Helseth, A., M. Fodstad, M. Askeland, B. Mo, O. Bjarte Nilsen, J. I. Pérez-Díaz, M. Chazarra, I. Guisández. 2017. Assessing hydropower operational profitability considering energy and reserve markets. *IET Renewable Power Generation* **11**(13) 1640–1647.
- Helseth, A., M. Fodstad, B. Mo. 2016. Optimal medium-term hydropower scheduling considering energy and reserve capacity markets. *IEEE Transactions on Sustainable Energy* **7**(3) 934–942.
- Helseth, A., B. Mo, M. Fodstad, M. Hjelmeland. 2015. Co-optimizing sales of energy and capacity in a hydropower scheduling model. *Proceedings of IEEE Eindhoven PowerTech*. 1–6.
- IEA. 2020a. Renewables 2020 data explorer. <https://www.iea.org/articles/renewables-2020-data-explorer?mode=market®ion=World&product=Hydro>. Accessed on 1st March 2021.
- IEA. 2020b. Renewables 2020 fuel report. <https://www.iea.org/reports/renewables-2020/hydropower-bioenergy-csp-and-geothermal#hydropower>. Accessed on 1st March 2021.
- Klæboe, G., J. Braathen, A. Eriksrud, S.-E. Fleten. 2019. Day-ahead market bidding taking the balancing power market into account. *Available on SSRN*.
- Löhndorf, N., D. Wozabal, S. Minner. 2013. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research* **61**(4) 810–823.
- Lorca, A., X. A. Sun, E. Litvinov, T. Zheng. 2016. Multistage adaptive robust optimization for the unit commitment problem. *Operations Research* **64**(1) 32–51.
- Mayer, J. 2014. Electricity production and spot-prices in Germany 2014. Fraunhofer Institute for Solar Energy Systems.
- Morris, C., M. Pehnt. 2015. Energy transition: The German Energiewende. Heinrich Böll Foundation.

- Pan, L., M. Housh, P. Liu, X. Cai, X. Chen. 2015. Robust stochastic optimization for reservoir operation. *Water Resources Research* **51**(1) 409–429.
- Pereira, M., L. Pinto. 1991. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* **52** 359–375.
- Philpott, A., V. de Matos. 2012. Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research* **218**(2) 470–483.
- Philpott, A., V. de Matos, L. Kapelevich. 2018. Distributionally robust SDDP. *Computational Management Science* **15** 431–454.
- Press, W., S. Teukolsky, W. Vetterling, B. Flannery. 2007. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press.
- Pritchard, G., A. Philpott, P. Neame. 2005. Hydroelectric reservoir optimization in a pool market. *Mathematical Programming* **103**(3) 445–461.
- Rockafellar, R. T., Roger J.-B. Wets. 2010. *Variational Analysis*. Springer.
- Schillinger, M., H. Weigt, M. Barry, R. Schumann. 2017. Hydropower operation in a changing environment – a Swiss case study. *Proceedings of the 14th International Conference on the European Energy Market*. 1–6.
- Shapiro, A., D. Dentcheva, A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM.
- Shapiro, A., W. Tekaya, M. Soares, J. da Costa. 2013. Worst-case-expectation approach to optimization under uncertainty. *Operations Research* **61**(6) 1435–1449.
- Vayanos, P., D. Kuhn, B. Rustem. 2012. A constraint sampling approach for multi-stage robust optimization. *Automatica* **48**(3) 459–471.
- Wirth, H. 2016. Recent facts about photovoltaics in Germany. Fraunhofer Institute for Solar Energy Systems.
- Yanikoğlu, I., B. L. Gorissen, D. den Hertog. 2019. A survey of adjustable robust optimization. *European Journal of Operational Research* **277**(3) 799–813.

Appendix.

A. Serial Independence of Reserve Activations

Throughout the paper we assume that the uncertain reserve activations ρ_t^u and ρ_t^v are serially independent across different hours $t \in \mathcal{T}$. To further investigate this assumption, recall that secondary reserve is dispatched by the grid operator in five-second intervals. It is therefore expedient to model the reserve activations as hourly averages of more fine-grained reserve activation processes $\rho^u(s)$ and $\rho^v(s)$, where $s \in \mathcal{S}$ indexes all five-second intervals within the planning horizon. Formally, we thus set

$$\rho_t^u = \frac{1}{|\mathcal{S}(t)|} \sum_{s \in \mathcal{S}(t)} \rho^u(s) \quad \text{and} \quad \rho_t^v = \frac{1}{|\mathcal{S}(t)|} \sum_{s \in \mathcal{S}(t)} \rho^v(s) \quad \forall t \in \mathcal{T}, \quad (7)$$

where $\mathcal{S}(t)$ denotes the set of all 720 five-second intervals in hour t . In any given interval $s \in \mathcal{S}$, there cannot be a simultaneous call-off on the reserve-up and reserve-down markets. This prompts us to model $\rho^u(s)$ and $\rho^v(s)$ as Bernoulli random variables, which evaluate to 0 if there is no call-off and to 1 if there is a call-off on the reserve-up or the reserve-down market, respectively, and to assume that $\rho^u(s)$ and $\rho^v(s)$ cannot evaluate to 1 at the same time. More precisely, we model the stochastic process $\{(\rho^u(s), \rho^v(s))\}_{s \in \mathcal{S}}$ as a time-homogeneous Markov chain with three states: no call-off (state 1: $\rho^u(s) = \rho^v(s) = 0$), a call-off on the reserve-up market (state 2: $\rho^u(s) = 1$ and $\rho^v(s) = 0$), or a call-off on the reserve-down market (state 3: $\rho^u(s) = 0$ and $\rho^v(s) = 1$). If we assume that the system cannot transition from a reserve-up call-off to a reserve-down call-off (or vice versa) in just five seconds and that call-offs on the reserve-up and the reserve-down markets are equally likely, then the transition probability matrix of the Markov chain can be represented as

$$\mathbf{P} = \begin{bmatrix} 1-2p & p & p \\ q & 1-q & 0 \\ q & 0 & 1-q \end{bmatrix},$$

where the rows and columns refer to the three states in order of appearance. Here, $p \in (0, \frac{1}{2})$ denotes the probability of a call-off on the reserve-up or the reserve-down market in the next five seconds if there is no call-off at present. Similarly, $q \in (0, 1)$ denotes the probability that an ongoing call-off on either market terminates over the next five seconds. The imposed time homogeneity makes the simplifying assumption that the probability of entering a new reserve activation (*i.e.*, the probability p leaving state 1) as well as the probability of terminating a reserve activation (*i.e.*, the probability q of entering state 1) are independent of the recent history of reserve activations. Using this Markov chain, we will investigate below to which degree the averaged hourly reserve activations ρ_t^u and ρ_t^v can be considered serially independent. To estimate p and q , we first note that the unique invariant distribution of the Markov chain encoded by \mathbf{P} is given by $\boldsymbol{\pi} = \frac{1}{q+2p}(q, p, p)$. Thus, the long-run fraction of time spent in reserve-up call-offs (or equivalently, reserve-down call-offs) is given by $\varrho = \frac{p}{q+2p}$. In addition, an elementary calculation reveals that $1/q$ coincides with the expected duration h of a call-off. Both ϱ and h can be estimated from data, and they uniquely determine p and q . In the following we assume that the initial state distribution is given by $\boldsymbol{\pi}$. Thus, the Markov chain is in the stationary regime, and $\mathbb{E}[\rho^u(s)] = \mathbb{E}[\rho^v(s)] = \varrho$ for all $s \in \mathcal{S}$, which in turn implies via (7) that $\hat{\rho}_t^u = \mathbb{E}[\rho_t^u] = \varrho$ and $\hat{\rho}_t^v = \mathbb{E}[\rho_t^v] = \varrho$ for all $t \in \mathcal{T}$. In this case, the dependence between the states of the Markov chain at time 1 and at time s can be measured by the uncertainty coefficient $U(s) =$




Figure 7 The uncertainty coefficient $U(s)$ as a function of time s measured in five-second intervals for different expected call durations h when the frequency of reserve-up call-offs equals $\varrho = 1\%$.

	$h = 1$ min	$h = 5$ mins	$h = 10$ mins	$h = 15$ mins	$h = 20$ mins
$t = 2$	0.01%	0.18%	1.26%	2.99%	5.24%
$t = 3$	$7 \times 10^{-3}\%$	0.03%	0.06%	0.08%	0.11%
$t = 4$	$7 \times 10^{-3}\%$	0.02%	0.06%	0.07%	0.08%

Table 3 The uncertainty coefficient between (ρ_1^u, ρ_1^v) and (ρ_t^u, ρ_t^v) for different hourly time lags t and expected call durations h when the frequency of reserve-up call-offs equals $\varrho = 1\%$.

$I((\rho^u(s), \rho^u(s)); (\rho^u(1), \rho^u(1))) / H((\rho^u(s), \rho^u(s)))$, where $I(\cdot; \cdot)$ denotes mutual information, and $H(\cdot)$ denotes entropy. The uncertainty coefficient ranges over $[0, 1]$ and characterizes the fraction of the bits of $(\rho^u(s), \rho^u(s))$ that can be predicted from $(\rho^u(1), \rho^u(1))$. The smaller $U(s)$, the more *independent* are the reserve activations at times 1 and s ; see *e.g.* Press et al. (2007, §14.7.4). Figure 7 shows that $U(s)$ decays rapidly with s for $\varrho = 1\%$ and for all values of $h \in \{1 \text{ min}, 5 \text{ min}, 10 \text{ min}, 15 \text{ min}, 20 \text{ min}\}$, which implies that the reserve activations become increasingly independent as the lag s grows. In particular, for lags exceeding one hour ($s \geq 720$), the uncertainty coefficient $U(s)$ virtually vanishes for all considered values of h . Note that 15 minutes is a natural upper bound on h as longer-lasting imbalances trigger an activation of tertiary reserves.

Table 3 reports the uncertainty coefficient between the hourly reserve activations (ρ_1^u, ρ_1^v) and (ρ_t^u, ρ_t^v) constructed as in (7) for $t \in \{2, 3, 4\}$. Here, the uncertainty coefficient is computed via simulation using 10^6 sample trajectories. Hence, consecutive hourly reserve activations are indeed nearly independent.

B. Proofs and Auxiliary Technical Results

Proof of Proposition 1 As problem (2) is a relaxation of problem (1), it suffices to show that any feasible solution of problem (2) corresponds to a feasible solution of problem (1) that attains the same objective function value. To this end, select any $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t, \mathbf{w}_d)\}_{t,d}$ feasible in (2). As spillage is neither restricted nor penalized and as the reservoir topology represents a directed acyclic graph, any amount of water equal to the slack of the reformulated reservoir balance constraints in (2) can be spilled through to node R sequentially for all reservoirs in the partial order induced by the topology graph and sequentially for all hours in the planning horizon. Thus, one can systematically replace the spillage decision \mathbf{z}_t with another $\mathcal{F}_{[t]}$ -measurable random vector \mathbf{z}'_t for each $t \in \mathcal{T}$ such that $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}'_t, \mathbf{w}_d)\}_{t,d}$ is feasible in problem (1). The claim now follows because the spillage decisions do not enter the objective function. \square

The proof of Proposition 2 relies on an auxiliary result which we state and prove first. To this end, consider a generic multistage stochastic program of the form

$$\begin{aligned} \sup \quad & \mathbb{E} [f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]})] \\ \text{s.t.} \quad & \mathbf{x}_0 \in \mathcal{L}^{n_0}(\mathcal{F}_0), \mathbf{x}_t \in \mathcal{L}^{n_t}(\mathcal{F}_{[t]}) \quad \forall t = 1, \dots, T, \end{aligned} \tag{8}$$

where the expectation is taken with respect to \mathbb{P} , the joint distribution of the random vectors $\boldsymbol{\xi}_t \in \mathbb{R}^{k_t}$, $t = 0, \dots, T$. Here, f constitutes a normal integrand in the sense of Rockafellar and Wets (2010, Definition 14.27), that is, its epigraph $\text{epi}(f(\cdot, \boldsymbol{\xi}_{[T]}))$ is a closed-valued measurable multifunction of $\boldsymbol{\xi}_{[T]}$. We assume that f may adopt the value $-\infty$ and thereby encode \mathbb{P} -almost sure constraints on the decision variables. We also retain the notational conventions for combined random vectors and σ -algebras introduced in Section 2. The following lemma relies on Theorem 14.60 by Rockafellar and Wets (2010) and the properties of conditional expectations.

LEMMA 1. *The optimal value of the multistage stochastic program (8) equals*

$$\mathbb{E} \left[\begin{array}{l} \sup \quad \mathbb{E} [f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \mid \boldsymbol{\xi}_0] \\ \text{s.t.} \quad \mathbf{x}_0 \in \mathbb{R}^{n_0}, \mathbf{x}_t \in \mathcal{L}^{n_t}(\mathcal{F}_{[1,t]}) \quad \forall t = 1, \dots, T \end{array} \right]. \tag{9}$$

Proof As joint optimization over the stage-wise decision variables is equivalent to sequential optimization, we can rewrite the multistage stochastic program (8) as

$$\sup_{\mathbf{x}_0 \in \mathcal{L}^{n_0}(\mathcal{F}_0)} \sup_{\mathbf{x}_1 \in \mathcal{L}^{n_1}(\mathcal{F}_{[1]})} \cdots \sup_{\mathbf{x}_{T-1} \in \mathcal{L}^{n_{T-1}}(\mathcal{F}_{[T-1]})} \sup_{\mathbf{x}_T \in \mathcal{L}^{n_T}(\mathcal{F}_{[T]})} \mathbb{E} [f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]})].$$

By Rockafellar and Wets (2010, Theorem 14.60), the above problem is equivalent to

$$\sup_{\mathbf{x}_0 \in \mathcal{L}^{n_0}(\mathcal{F}_0)} \sup_{\mathbf{x}_1 \in \mathcal{L}^{n_1}(\mathcal{F}_{[1]})} \cdots \sup_{\mathbf{x}_{T-1} \in \mathcal{L}^{n_{T-1}}(\mathcal{F}_{[T-1]})} \mathbb{E} \left[\sup_{\mathbf{x}_T \in \mathbb{R}^{n_T}} f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \right].$$

Recalling the tower property of conditional expectations, we can rewrite the resulting problem as

$$\sup_{\mathbf{x}_0 \in \mathcal{L}^{n_0}(\mathcal{F}_0)} \sup_{\mathbf{x}_1 \in \mathcal{L}^{n_1}(\mathcal{F}_{[1]})} \cdots \sup_{\mathbf{x}_{T-1} \in \mathcal{L}^{n_{T-1}}(\mathcal{F}_{[T-1]})} \mathbb{E} \left[\mathbb{E} \left[\sup_{\mathbf{x}_T \in \mathbb{R}^{n_T}} f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \mid \boldsymbol{\xi}_{[T-1]} \right] \right]$$

and invoke the interchangeability theorem of Rockafellar and Wets (2010) once again to prove its equivalence to

$$\sup_{\mathbf{x}_0 \in \mathcal{L}^{n_0}(\mathcal{F}_0)} \sup_{\mathbf{x}_1 \in \mathcal{L}^{n_1}(\mathcal{F}_{[1]})} \cdots \mathbb{E} \left[\sup_{\mathbf{x}_{T-1} \in \mathbb{R}^{n_{T-1}}} \mathbb{E} \left[\sup_{\mathbf{x}_T \in \mathbb{R}^{n_T}} f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \mid \boldsymbol{\xi}_{[T-1]} \right] \right]$$

Repeating this argument, it is possible to move all supremum operators inside the appropriate conditional expectation layers such as to obtain the equivalent problem

$$\mathbb{E} \left[\sup_{\mathbf{x}_0 \in \mathbb{R}^{n_0}} \mathbb{E} \left[\sup_{\mathbf{x}_1 \in \mathbb{R}^{n_1}} \mathbb{E} \left[\sup_{\mathbf{x}_2 \in \mathbb{R}^{n_2}} \cdots \mathbb{E} \left[\sup_{\mathbf{x}_T \in \mathbb{R}^{n_T}} f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \mid \boldsymbol{\xi}_{[T-1]} \right] \cdots \mid \boldsymbol{\xi}_{[1]} \mid \boldsymbol{\xi}_0 \right] \right] \right].$$

Applying the interchangeability theorem of Rockafellar and Wets (2010) in the reverse direction, we can then move the second supremum out of its conditional expectation layer to obtain

$$\mathbb{E} \left[\sup_{\mathbf{x}_0 \in \mathbb{R}^{n_0}} \sup_{\mathbf{x}_1 \in \mathcal{L}^{n_1}(\mathcal{F}_{[1,1]})} \mathbb{E} \left[\sup_{\mathbf{x}_2 \in \mathbb{R}^{n_2}} \cdots \mathbb{E} \left[\sup_{\mathbf{x}_T \in \mathbb{R}^{n_T}} f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \mid \boldsymbol{\xi}_{[T-1]} \right] \cdots \mid \boldsymbol{\xi}_0 \right] \right].$$

Iterating this argument from the outside to the inside, it is indeed possible to move all supremum operators into the outermost (unconditional) expectation layer and thus obtain

$$\mathbb{E} \left[\sup_{\mathbf{x}_0 \in \mathbb{R}^{n_0}} \sup_{\mathbf{x}_1 \in \mathcal{L}^{n_1}(\mathcal{F}_{[1,1]})} \sup_{\mathbf{x}_2 \in \mathcal{L}^{n_2}(\mathcal{F}_{[1,2]})} \cdots \sup_{\mathbf{x}_T \in \mathcal{L}^{n_T}(\mathcal{F}_{[1,T]})} \mathbb{E} \left[f(\mathbf{x}_0, \dots, \mathbf{x}_T, \boldsymbol{\xi}_{[T]}) \mid \boldsymbol{\xi}_0 \right] \right].$$

Re-combining the stage-wise supremum operators finally yields problem (9). \square

Proof of Proposition 2 Observe that if the reservoir filling levels $\{\mathbf{w}_d\}_{d \in \mathcal{D}}$ are fixed, then the remaining decisions in problem (2) are no longer coupled across days. This allows us to decompose the stochastic program (2) into an outer maximization problem over the end-of-day reservoir filling levels $\{\mathbf{w}_d\}_{d \in \mathcal{D}}$ and a series of (mutually independent) inner stochastic programs maximizing over the bidding decisions $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(d)}$ and the operational decisions $\{(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$, one for each day $d \in \mathcal{D}$. Formally, problem (2) is thus equivalent to the outer stochastic program

$$\begin{aligned} & \sup \sum_{d \in \mathcal{D}} \hat{\Pi}_d(\mathbf{w}_{d-1}, \mathbf{w}_d) \\ & \text{s.t. } \mathbf{w}_d \in \mathcal{L}^R(\mathcal{F}_{[\uparrow(d)]}) \quad \forall d \in \mathcal{D} \\ & \quad \underline{\mathbf{w}}_d \leq \mathbf{w}_d \leq \bar{\mathbf{w}}_d \quad \forall d \in \mathcal{D}, \mathbb{P}\text{-a.s.}, \end{aligned}$$

where $\hat{\Pi}_d(\mathbf{w}_{d-1}, \mathbf{w}_d)$ stands for the optimal value of the inner (parametric) stochastic program

$$\begin{aligned} & \sup \mathbb{E} \left[\sum_{t \in \mathcal{T}(d)} \pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t \right] \\ & \text{s.t. } s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[\downarrow(d)]}), \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[t]}) \quad \forall t \in \mathcal{T}(d) \\ & \quad 0 \leq u_t, 0 \leq v_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\ & \quad s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\ & \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)} \phi_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}\text{-a.s.} \\ & \quad \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \phi_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \mathbb{P}\text{-a.s.}, \end{aligned} \tag{10}$$

which optimizes only over decisions pertaining to day d . Observe that problem (10) differs from the trader's problem (CT) in that it maximizes an *unconditional* expectation of the profits earned on day d and in that its decision variables may still adapt to information revealed *prior* to day d .

We will now use Lemma 1 to show that $\hat{\Pi}_d(\mathbf{w}_{d-1}, \mathbf{w}_d)$ coincides with the expected value of $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})$. For notational convenience, we abbreviate the random variables revealed before the

beginning of day d as $\boldsymbol{\xi}_0^d = \boldsymbol{\xi}_{[\downarrow(d)]}$ and those revealed up to each hour $h = 1, \dots, H$ of day d by $\boldsymbol{\xi}_{[h]}^d = \boldsymbol{\xi}_{[\downarrow(d)+h]}$. Furthermore, we gather the hourly decisions of day d in the vectors

$$\begin{aligned}\mathbf{x}_0^d &= (\{s_\tau, u_\tau, v_\tau\}_{\tau=\downarrow(d)}^{\uparrow(d)}) \in \mathcal{L}^{3H}(\mathcal{F}_{[\downarrow(d)]}), \\ \mathbf{x}_h^d &= (\mathbf{g}_{\downarrow(d)+h}, \mathbf{p}_{\downarrow(d)+h}, \mathbf{z}_{\downarrow(d)+h}) \in \mathcal{L}^{3A}(\mathcal{F}_{[\downarrow(d)+h]}) \quad \forall h = 1, \dots, H.\end{aligned}$$

Using this notation, the profit earned on day d can be expressed more concisely as

$$f^d(\mathbf{x}_0^d, \dots, \mathbf{x}_H^d, \boldsymbol{\xi}_{[H]}^d) = \begin{cases} \sum_{t \in \mathcal{T}(d)} \pi_t^s s_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) u_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) v_t & \text{if } (\mathbf{x}_0^d, \dots, \mathbf{x}_H^d) \text{ satisfies all } \mathbb{P}\text{-almost sure} \\ & \text{constraints of problem (10) in scenario } \boldsymbol{\xi}_{[H]}^d, \\ -\infty & \text{otherwise.} \end{cases}$$

Problem (10) can thus be represented abstractly as

$$\hat{\Pi}_d(\mathbf{w}_{d-1}, \mathbf{w}_d) = \begin{cases} \sup \mathbb{E}[f^d(\mathbf{x}_0^d, \dots, \mathbf{x}_H^d, \boldsymbol{\xi}_{[H]}^d)] \\ \text{s.t. } \mathbf{x}_0^d \in \mathcal{L}^{3H}(\mathcal{F}_{[\downarrow(d)]}), \mathbf{x}_h^d \in \mathcal{L}^{3A}(\mathcal{F}_{[\downarrow(d)+h]}) \quad \forall h = 1, \dots, H. \end{cases}$$

By the extended interchangeability principle established in Lemma 1, we thus find

$$\hat{\Pi}_d(\mathbf{w}_{d-1}, \mathbf{w}_d) = \mathbb{E} \left[\begin{array}{l} \sup \mathbb{E}[f^d(\mathbf{x}_0^d, \dots, \mathbf{x}_H^d, \boldsymbol{\xi}_{[H]}^d) \mid \boldsymbol{\xi}_0^d] \\ \text{s.t. } \mathbf{x}_0^d \in \mathbb{R}^{3H}, \mathbf{x}_h^d \in \mathcal{L}^{3A}(\mathcal{F}_{[\downarrow(d), \downarrow(d)+h]}) \quad \forall h = 1, \dots, H \end{array} \right].$$

Unravelling the abbreviations shows that the stochastic program inside the expectation coincides with the trader's problem (CT), whose optimal value is given by $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\downarrow(d)]})$. \square

Proof of Proposition 3 Throughout this proof we will fix a realization of the contextual covariates $\boldsymbol{\xi}_{[\downarrow(d)]}$, which reflect the trader's information when solving problem (CT) for day d . Approximation 1 then implies that \mathbf{w}_{d-1} and \mathbf{w}_d reduce to deterministic constants. In addition, $\mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}$ becomes an *unconditional* distribution. Below we will denote the expectation with respect to this distribution by $\mathbb{E}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}[\cdot]$.

As problem (3) is a restriction of problem (CT), it suffices to show that for every feasible solution $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ of (CT) there is a feasible solution $\{(s'_t, u'_t, v'_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ of (12) that attains the same objective value. Such a solution can readily be constructed as

$$(s'_t, u'_t, v'_t) = (s_t, u_t, v_t) \quad \text{and} \quad (\mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t) = \mathbb{E}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}[(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t) \mid \mathcal{F}_{[\downarrow(d), t]}^p] \quad \forall t \in \mathcal{T}(d).$$

The equality of objective values is immediate because the objective function only depends on the bidding decisions, which are preserved. The non-anticipativity constraints $\mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]}^p)$ are also satisfied thanks to the defining properties of conditional expectations. To show that the almost sure constraints hold, we recall first that the reserve activations are serially independent and independent of all other sources of uncertainty. This implies that $\{(\rho_\tau^u, \rho_\tau^v)\}_{\tau=t+1}^{\uparrow(d)}$ is independent of the non-anticipative flow decisions $\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]})$ under the distribution $\mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}$, that is

$$(\mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t) = \mathbb{E}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}[(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t) \mid \mathcal{F}_{[\downarrow(d), \uparrow(d)]}^p] \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \quad (11)$$

for all $t \in \mathcal{T}(d)$. The feasibility of $\{(s'_t, u'_t, v'_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ in (3) therefore follows from the feasibility of $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ in (CT) and the linearity of the almost sure constraints in (3), which implies

that they all remain valid under conditional expectations. For example, the newly constructed solution $\{(s'_t, u'_t, v'_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ obeys the end-of-day reservoir bound in (12) because

$$\begin{aligned} \mathbf{w}_d &= \mathbb{E}_{|\boldsymbol{\xi}_{[\Psi(d)]}}[\mathbf{w}_d | \mathcal{F}_{[\downarrow(d), \uparrow(d)]}^\rho] \\ &\leq \mathbb{E}_{|\boldsymbol{\xi}_{[\Psi(d)]}}[\mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) | \mathcal{F}_{[\downarrow(d), \uparrow(d)]}^\rho] \\ &= \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}'_\tau - \mathbf{p}'_\tau + \mathbf{z}'_\tau) \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}} \text{ a.s.}, \end{aligned}$$

where the first equality follows from Approximation 1, whereby \mathbf{w}_d is deterministic when $\boldsymbol{\xi}_{[\Psi(d)]}$ is kept fixed, the inequality holds because $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ satisfies the last constraint of (CT), and the last equality follows from (11). Thus, the claim follows. \square

Proof of Proposition 4 Any feasible solution of the stochastic program (3) satisfies the energy delivery constraint $s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t$ for all $t \in \mathcal{T}(d)$ $\mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}$ -almost surely. Similarly, for any $t \in \mathcal{T}(d)$, we have that $\mathbf{g}_t \geq \mathbf{0}$ and $\mathbf{p}_t \leq \bar{\mathbf{p}}_t$ $\mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}$ -almost surely. Together, these inequalities imply

$$s_t + \kappa_t^u \rho_t^u u_t - \kappa_t^v \rho_t^v v_t = \boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t \geq -\boldsymbol{\zeta}_t^\top \mathbf{p}_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}} \text{ a.s.}$$

As the scenario $(0, 1, 0, \bar{\rho}_t^v)$ belongs to the support of the random vector $(\kappa_t^u, \kappa_t^v, \rho_t^u, \rho_t^v)$ under the conditional probability measure $\mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}$, we thus find $s_t - \bar{\rho}_t^v v_t \geq -\boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t$ for every $t \in \mathcal{T}(d)$. \square

Proof of Proposition 5 We will prove the proposition by showing that every feasible solution of problem (4) corresponds to a feasible solution of problem (CT^r) with the same objective value. To this end, fix any feasible solution $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ of problem (4) and an arbitrary realization of the contextual covariates $\boldsymbol{\xi}_{[\Psi(d)]}$. Since all non-dummy reservoirs and pumps have finite capacity, there is $\bar{\mathbf{z}} \in \mathbb{R}^A$ such that $\mathbf{z}_t \leq \bar{\mathbf{z}}$ $\mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}$ -almost surely for all $t \in \mathcal{T}(d)$. Then, there exists an event $\Omega_0 \in \mathcal{F}$ with $\mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}[\Omega_0] = 1$ such that all constraints of problem (4) are satisfied pointwise for all $\omega \in \Omega_0$. In addition, by our assumptions about the statistics of the reserve allocations and activations, there exists a sequence of samples $\omega_n \in \Omega_0$, $n \in \mathbb{N}$, such that

$$\lim_{n \rightarrow \infty} \kappa_t^u(\omega_n) = 1, \quad \lim_{n \rightarrow \infty} \rho_t^u(\omega_n) = \bar{\rho}_t^u \quad \text{and} \quad \lim_{n \rightarrow \infty} \rho_t^v(\omega_n) = 0 \quad \forall t \in \mathcal{T}(d).$$

Note that the sequence $\{(\mathbf{g}_t(\omega_n), \mathbf{p}_t(\omega_n), \mathbf{z}_t(\omega_n))\}_{n \in \mathbb{N}}$ is bounded. By passing to a sub-sequence if necessary, we may thus assume that this sequence converges for every $t \in \mathcal{T}(d)$, and we may define

$$\mathbf{g}'_t = \lim_{n \rightarrow \infty} \mathbf{g}_t(\omega_n), \quad \mathbf{p}'_t = \lim_{n \rightarrow \infty} \mathbf{p}_t(\omega_n) \quad \text{and} \quad \mathbf{z}'_t = \lim_{n \rightarrow \infty} \mathbf{z}_t(\omega_n) \quad \forall t \in \mathcal{T}(d).$$

One readily verifies that $\{(s_t, u_t, v_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ is feasible in (CT^r) because the constraints of the stochastic program (4) are continuous in the uncertain parameters and the uncertain flow decisions. Also, the objective value of $\{(s_t, u_t, v_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ in (CT^r) coincides with that of $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ in (4) because the two solutions involve identical market bids. The claim thus follows. \square

Proof of Theorem 1 This is an immediate consequence of Propositions 3, 4 and 5, as well as Proposition 9 from Appendix C. \square

Proof of Theorem 2 From Propositions 1 and 2 we know that the bidding model (C) has the same optimal value as the planner's problem (CP), and it is easy to see that problem (CP^r) is obtained by applying Approximation 1 to problem (CP). As this approximation consists in restricting the planner's information structure, the optimal value of problem (CP^r) is no larger than that of problem (C). It remains to be shown that (CP^r) and (C^r) share the same optimal value. As Approximation 1 is in force, we may conclude via Theorem 1 that the optimal value $\Pi_d(\mathbf{w}_{d-1}, \mathbf{w}_d, \boldsymbol{\xi}_{[\psi(d)]})$ of the stochastic program (CT) coincides with the optimal value of the linear program (CT^r). Substituting (CT^r) into (CP^r) and using Theorem 14.60 by Rockafellar and Wets (2010) to move the maximization over the trading and flow decisions out of the expectation and the sum finally yields (C^r). \square

C. Proof that Problem (CT) is a Restriction of Problem (CT^r)

In this appendix we prove that problem (CT) is a restriction of problem (CT^r), that is, that the arc from the node labelled (CT^r) to the node labelled (CT) in Figure 1 is justified. To this end, we introduce a variant of problem (CT) that we henceforth refer to as the *individual* trader's problem.

$$\begin{aligned}
& \sup \quad \mathbb{E} \left[\sum_{t \in \mathcal{T}(d)} \pi_t^s \mathbf{1}^\top \mathbf{s}_t + \kappa_t^u (\pi_t^u + \rho_t^u \psi_t^u) \mathbf{1}^\top \mathbf{u}_t + \kappa_t^v (\pi_t^v + \rho_t^v \psi_t^v) \mathbf{1}^\top \mathbf{v}_t \mid \boldsymbol{\xi}_{[\downarrow(d)]} \right] \\
& \text{s.t.} \quad \mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t \in \mathbb{R}^A, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]}) \quad \forall t \in \mathcal{T}(d) \\
& \quad \mathbf{0} \leq \mathbf{u}_t, \mathbf{0} \leq \mathbf{v}_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\
& \quad \mathbf{s}_t + \kappa_t^u \rho_t^u \mathbf{u}_t - \kappa_t^v \rho_t^v \mathbf{v}_t = \boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \quad (\text{IT}) \\
& \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.} \\
& \quad \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\downarrow(d)]}}\text{-a.s.}
\end{aligned}$$

To understand how (IT) differs from (CT), recall that energy is produced by the generators and consumed by the pumps installed along the arcs of the reservoir topology. Under problem (IT), the company places an individual bid for every arc of its reservoir system. We gather these individual (arc-specific) bids placed on the spot market, the reserve-up market and the reserve-down market in the vectors $\mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(t)])}$ [MWh], respectively. In contrast, the previous trader's problem (CT) contains aggregate decisions $s_t, u_t, v_t \in \mathcal{L}(\mathcal{F}_{[\downarrow(t)])}$. Note that (CT) is non-inferior to (IT) because it increases the producer's flexibility in choosing the generation and pumping decisions. We show below that, similarly to (CT), problem (IT) can undergo several simplifications. However, we will show that (IT) itself eventually reduces to the tractable linear program

$$\begin{aligned}
& \sup \quad \sum_{t \in \mathcal{T}(d)} \tilde{\pi}_t^s \mathbf{1}^\top \mathbf{s}_t + \tilde{\pi}_t^u \mathbf{1}^\top \mathbf{u}_t + \tilde{\pi}_t^v \mathbf{1}^\top \mathbf{v}_t \\
& \text{s.t.} \quad \mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t \in \mathbb{R}^A, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathbb{R}^A \quad \forall t \in \mathcal{T}(d) \\
& \quad \mathbf{0} \leq \mathbf{u}_t, \mathbf{0} \leq \mathbf{v}_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}(d) \\
& \quad \mathbf{s}_t + \bar{\rho}_t^u \mathbf{u}_t = \boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t \quad \forall t \in \mathcal{T}(d) \\
& \quad \mathbf{s}_t - \bar{\rho}_t^v \mathbf{v}_t \geq -\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t \quad \forall t \in \mathcal{T}(d) \\
& \quad \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(d) \\
& \quad \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau)
\end{aligned} \quad (\text{IT}^r)$$

if we impose Approximation 1, whereby \mathbf{w}_d is restricted to $\mathcal{L}^R(\mathcal{F}_{[\uparrow(d-1)])}$. By construction, problem (CT^r) is non-inferior to (IT^r). Proposition 9 below, which has previously been used to prove Theorem 1, shows that (IT^r) is non-inferior to (CT^r) as well.

In analogy to Proposition 3, we first establish that, conditional on $\boldsymbol{\xi}_{[\downarrow(d)]}$, the wait-and-see decisions $\{(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ in (IT) may be restricted to measurable functions of the reserve allocations and the reserve

activations without sacrificing optimality. This additional information restriction gives rise to the problem

$$\begin{aligned}
& \sup \sum_{t \in \mathcal{T}(d)} \tilde{\pi}_t^s \mathbf{1}^\top \mathbf{s}_t + \tilde{\pi}_t^u \mathbf{1}^\top \mathbf{u}_t + \tilde{\pi}_t^v \mathbf{1}^\top \mathbf{v}_t \\
& \text{s.t. } \mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t \in \mathbb{R}^A, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]}^\rho) \quad \forall t \in \mathcal{T}(d) \\
& \mathbf{0} \leq \mathbf{u}_t, \mathbf{0} \leq \mathbf{v}_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& \mathbf{s}_t + \kappa_t^u \rho_t^u \mathbf{u}_t - \kappa_t^v \rho_t^v \mathbf{v}_t = \boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.},
\end{aligned} \tag{12}$$

PROPOSITION 6. *Under Approximation 1, the optimal values of problems (IT) and (12) are equal.*

Proof As in the proof of Proposition 3, one may condition any feasible solution of problem (IT) on the history of reserve activations to construct a feasible solution of problem (12) that adopts the same objective value. Details are omitted for brevity. \square

Proposition 6 asserts that the individual trader's problem (IT) is equivalent to the stochastic program (12), which accommodates only $H + 1$ decision stages and whose wait-and-see decisions depend only on the reserve allocations and the reserve activations. However, problem (12) still constitutes an infinite-dimensional linear program. We now show that problem (12) is indeed equivalent to an efficiently solvable finite linear program. As a first step towards this goal, we derive a family of valid inequalities that may be added to problem (12) without affecting its feasible set.

PROPOSITION 7. *Any feasible solution of (12) satisfies $\mathbf{s}_t - \bar{\rho}_t^v \mathbf{v}_t \geq -\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t$ for all $t \in \mathcal{T}(d)$.*

Proof The proof widely parallels that of Proposition 4 and is thus omitted for brevity. \square

By Proposition 7, the deterministic inequalities $\mathbf{s}_t - \bar{\rho}_t^v \mathbf{v}_t \geq -\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t$ are valid for all $t \in \mathcal{T}(d)$ and can therefore be appended to problem (12) without affecting its feasible set. Observe that $\mathbf{s}_t + \boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t$ comprises the arc-wise maximum bids in the reserve-down market on which the company can deliver in case of a call-off. To see this, note that $s_{t,a}$ is the amount of energy produced on arc a for the spot market and that $\zeta_{t,a} \bar{p}_{t,a}$ represents the maximum amount of energy that can be absorbed on arc a by pumping. In case of a call-off on the reserve-down market, the energy production on arc a can thus be reduced at most by $s_{t,a} + \zeta_{t,a} \bar{p}_{t,a}$.

Appending the valid inequalities to (12), we obtain

$$\begin{aligned}
& \sup \sum_{t \in \mathcal{T}(d)} \tilde{\pi}_t^s \mathbf{1}^\top \mathbf{s}_t + \tilde{\pi}_t^u \mathbf{1}^\top \mathbf{u}_t + \tilde{\pi}_t^v \mathbf{1}^\top \mathbf{v}_t \\
& \text{s.t. } \mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t \in \mathbb{R}^A, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t \in \mathcal{L}^A(\mathcal{F}_{[\downarrow(d), t]}^\rho) \quad \forall t \in \mathcal{T}(d) \\
& \mathbf{0} \leq \mathbf{u}_t, \mathbf{0} \leq \mathbf{v}_t, \mathbf{0} \leq \mathbf{g}_t \leq \bar{\mathbf{g}}_t, \mathbf{0} \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \mathbf{0} \leq \mathbf{z}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& \mathbf{s}_t + \kappa_t^u \rho_t^u \mathbf{u}_t - \kappa_t^v \rho_t^v \mathbf{v}_t = \boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& \mathbf{s}_t - \bar{\rho}_t^v \mathbf{v}_t \geq -\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t \quad \forall t \in \mathcal{T}(d) \\
& \underline{\mathbf{w}}_t \leq \mathbf{w}_{d-1} + \sum_{\tau=\downarrow(d)}^t \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \leq \bar{\mathbf{w}}_t \quad \forall t \in \mathcal{T}(d), \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.} \\
& \mathbf{w}_d \leq \mathbf{w}_{d-1} + \sum_{\tau \in \mathcal{T}(d)} \boldsymbol{\phi}_\tau + \mathbf{M}(\mathbf{g}_\tau - \mathbf{p}_\tau + \mathbf{z}_\tau) \quad \mathbb{P}_{|\xi_{[\downarrow(d)]}}\text{-a.s.},
\end{aligned} \tag{13}$$

which is equivalent to problem (12) by virtue of Proposition 7.

Next, we will show that problem (13) is equivalent to the reduced individual trader's problem (IT^r). Our argument critically relies on the following technical yet intuitive lemma, which asserts that there is no benefit in simultaneous generation and pumping.

LEMMA 2. *The optimal value of the reduced planner's problem (IT^r) does not decrease if we append the complementarity constraints $\mathbf{g}_t \circ \mathbf{p}_t = \mathbf{0}$ for all $t \in \mathcal{T}(d)$.*

Proof Consider any feasible solution $\{(\mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ of problem (IT^r), and construct a new solution $\{(\mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ with adjusted flow decisions

$$\mathbf{g}'_t = \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \boldsymbol{\Delta}_t, \quad \mathbf{p}'_t = \mathbf{p}_t - \boldsymbol{\eta}_t \circ \boldsymbol{\Delta}_t, \quad \mathbf{z}'_t = \mathbf{z}_t + (\boldsymbol{\zeta}_t - \boldsymbol{\eta}_t) \circ \boldsymbol{\Delta}_t,$$

where $\boldsymbol{\Delta}_t \in \mathbb{R}^A$ is defined through $\Delta_{t,a} = \min\{g_{t,a}/\zeta_{t,a}, p_{t,a}/\eta_{t,a}\}$. This new solution preserves the market decisions and consequently the objective value of the original solution, and it is readily seen to satisfy the complementarity constraints $\mathbf{g}'_t \circ \mathbf{p}'_t = \mathbf{0}$ for all $t \in \mathcal{T}(d)$. The claim thus follows if we can show that the new solution is feasible in (IT^r). To this end, note that $\boldsymbol{\Delta}_t \geq \mathbf{0}$, which implies that $\mathbf{g}'_t \leq \mathbf{g}_t \leq \bar{\mathbf{g}}$, $\mathbf{p}'_t \leq \mathbf{p}_t \leq \bar{\mathbf{p}}$, and $\mathbf{z}'_t \geq \mathbf{0}$, where the last inequality exploits our standing assumption that $\boldsymbol{\zeta}_t > \boldsymbol{\eta}_t$. Similarly, the inequalities $\boldsymbol{\zeta}_t \circ \boldsymbol{\Delta}_t \leq \mathbf{g}_t$ and $\boldsymbol{\eta}_t \circ \boldsymbol{\Delta}_t \leq \mathbf{p}_t$ ensure that $\mathbf{g}'_t \geq \mathbf{0}$ and $\mathbf{p}'_t \geq \mathbf{0}$. Finally, observe that $\mathbf{g}'_t - \mathbf{p}'_t + \mathbf{z}'_t = \mathbf{g}_t - \mathbf{p}_t + \mathbf{z}_t$ (*i.e.*, the net reservoir outflows remain unchanged) and $\boldsymbol{\eta}_t \circ \mathbf{g}'_t - \boldsymbol{\zeta}_t \circ \mathbf{p}'_t = \boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t$ (*i.e.*, the arc-wise net energy production quantities remain unchanged), which implies that the new solution satisfies the reservoir balance constraints and the energy delivery constraints, respectively. In summary, we have shown that the new solution is indeed feasible in (IT^r), and thus the claim follows. \square

PROPOSITION 8. *Under Approximation 1, the optimal values of (13) and (IT^r) are equal.*

Proof We will prove the proposition by showing that every feasible solution of problem (13) corresponds to a feasible solution of problem (IT^r) with the same objective value and vice versa. Since the first correspondence follows from a simple adaptation of the respective argument in the proof of Proposition 5, we omit its detailed proof for the sake of brevity.

To prove that every feasible solution of problem (IT^r) corresponds to a feasible solution of problem (13), we fix any feasible solution $\{(\mathbf{s}_t, \mathbf{u}_t, \mathbf{v}_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ of the reduced trader's problem (IT^r). By Lemma 2, we may assume without loss of generality that this solution satisfies the complementarity constraints $\mathbf{g}_t \circ \mathbf{p}_t = \mathbf{0}$ for all $t \in \mathcal{T}(d)$. We claim that one can systematically construct flow decisions $\{(\mathbf{g}'_t, \mathbf{p}'_t)\}_{t \in \mathcal{T}(d)}$ that satisfy

$$g'_{t,a}, p'_{t,a} \in \mathcal{L}(\mathcal{F}_{\downarrow(d),t}^{\rho}) \tag{14a}$$

$$0 \leq g'_{t,a} \leq \bar{g}_{t,a} \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}\text{-a.s.} \tag{14b}$$

$$0 \leq p'_{t,a} \leq \bar{p}_{t,a} \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}\text{-a.s.} \tag{14c}$$

$$\eta_{t,a} g'_{t,a} - \zeta_{t,a} p'_{t,a} = s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a} \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}\text{-a.s.} \tag{14d}$$

$$g'_{t,a} - p'_{t,a} \leq g_{t,a} - p_{t,a} \quad \mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}\text{-a.s.} \tag{14e}$$

for all $t \in \mathcal{T}(d)$ and $a \in \mathcal{A}$. As $g_{t,a} p_{t,a} = 0$, we can distinguish two cases for each hour-arc pair (t, a) , which necessitate two different constructions of the corresponding flows $g'_{t,a}$ and $p'_{t,a}$. Specifically, we first assume that $g_{t,a} = 0$ (Case 1), and then we assume that $g_{t,a} > 0$ (Case 2). From now on all equalities and inequalities involving random variables are understood to hold $\mathbb{P}_{|\boldsymbol{\xi}_{[\Psi(d)]}}$ -almost surely.

Case 1 ($g_{t,a} = 0$): We set $g'_{t,a} = 0$ and $p'_{t,a} = -(s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})/\zeta_{t,a}$. It is easy to verify that $(g'_{t,a}, p'_{t,a})$ satisfies (14b) and (14d). The non-anticipativity constraints (14a) are also met because the reserve allocations κ_t^u and κ_t^v are revealed in hour $\downarrow(t)$, whereas the reserve activations ρ_t^u and ρ_t^v are revealed in hour t . In order to establish (14c) and (14e), we first observe that the given feasible solution of the reduced trader's problem (IT^r) satisfies

$$s_{t,a} + \bar{\rho}_t^u u_{t,a} = \eta_{t,a} g_{t,a} - \zeta_{t,a} p_{t,a} \implies p_{t,a} = -(s_{t,a} + \bar{\rho}_t^u u_{t,a})/\zeta_{t,a},$$

where the implication holds because $g_{t,a} = 0$. As $\rho_t^u \leq \bar{\rho}_t^u$ and $\kappa_t^u \leq 1$ while ρ_t^v , κ_t^v , $u_{t,a}$ and $v_{t,a}$ are non-negative, we may thus conclude that $p_{t,a} \leq p'_{t,a}$. Hence, the constructed pumping decision $p'_{t,a}$ meets requirement (14e). Finally, requirement (14c) is satisfied because $0 \leq p_{t,a} \leq p'_{t,a}$ and because

$$p'_{t,a} = -(s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})/\zeta_{t,a} \leq -(s_{t,a} - \bar{\rho}_t^v v_{t,a})/\zeta_{t,a} \leq \bar{p}_{t,a},$$

where the second inequality follows from the valid cut derived in Proposition 7, which constitutes one of the constraints of the reduced trader's problem (IT^r). Thus, $g'_{t,a}$ and $p'_{t,a}$ satisfy (14).

Case 2 ($g_{t,a} > 0$): We set $g'_{t,a} = (s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})^+/\eta_{t,a}$ and $p'_{t,a} = (s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})^-/\zeta_{t,a}$. These flow decisions manifestly satisfy the non-anticipativity constraints (14a).

If $s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a} \leq 0$, then we have $g'_{t,a} = 0$ and $p'_{t,a} = -(s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})/\zeta_{t,a}$, and one can proceed as in Case 1 to show that the requirements (14b)–(14e) are met. From now on assume that $s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a} > 0$, in which case $g'_{t,a} = (s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})/\eta_{t,a}$ and $p'_{t,a} = 0$. Note first that the requirements (14c) and (14d) are trivially met. We further find that

$$g'_{t,a} = (s_{t,a} + \kappa_t^u \rho_t^u u_{t,a} - \kappa_t^v \rho_t^v v_{t,a})/\eta_{t,a} \leq (s_{t,a} + \bar{\rho}_t^u u_{t,a})/\eta_{t,a} = g_{t,a} \leq \bar{g}_{t,a}, \quad (15)$$

where the second equality follows from the constraint $s_{t,a} + \bar{\rho}_t^u u_{t,a} = \eta_{t,a} g_{t,a} - \zeta_{t,a} p_{t,a}$ of problem (IT^r) and the complementarity condition $g_{t,a} p_{t,a} = 0$, which implies that $p_{t,a} = 0$. Hence, requirement (14b) is satisfied. Finally, as $p'_{t,a} = p_{t,a} = 0$, the inequality (15) also implies that requirement (14e) is met. Thus, $g'_{t,a}$ and $p'_{t,a}$ satisfy again all of the requirements listed in (14).

Given the flow decisions $\{(g'_t, p'_t)\}_{t \in \mathcal{T}(d)}$ constructed above, we are now ready to introduce compatible spill decisions $z'_t = z_t + (g_t - p_t) - (g'_t - p'_t)$ for all $t \in \mathcal{T}(d)$. In the remainder of the proof we will demonstrate that the constructed solution $\{(s_t, u_t, v_t, g'_t, p'_t, z'_t)\}_{t \in \mathcal{T}(d)}$ is feasible in (13).

Note first that we need not be concerned with the constraints that only involve the market decisions $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(d)}$, which are trivially satisfied because $\{(s_t, u_t, v_t, g_t, p_t, z_t)\}_{t \in \mathcal{T}(d)}$ is feasible in problem (IT^r). Moreover, all constraints of problem (13) that do *not* involve the spill decisions are satisfied because of (14). It remains to verify that the spill decisions are non-anticipative as well as non-negative and that the reservoir balance constraints are satisfied. To this end, we note first that z'_t inherits non-anticipativity from g'_t and p'_t . Similarly, z'_t inherits non-negativity from z_t thanks to (14e). Finally, we highlight that (g'_t, p'_t, z'_t) impacts the reservoir balance constraints only through the net water outflows $g'_t - p'_t + z'_t$, which coincide with $g_t - p_t + z_t$ by the construction of z'_t . This guarantees via the feasibility of $\{(s_t, u_t, v_t, g_t, p_t, z_t)\}_{t \in \mathcal{T}(d)}$ in (IT^r) that the reservoir balance constraints are satisfied. Therefore, $\{(s_t, u_t, v_t, g'_t, p'_t, z'_t)\}_{t \in \mathcal{T}(d)}$ is indeed feasible in (13).

The claim now follows because $\{(s_t, \mathbf{u}_t, \mathbf{v}_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ and $\{(s_t, \mathbf{u}_t, \mathbf{v}_t, \mathbf{g}'_t, \mathbf{p}'_t, \mathbf{z}'_t)\}_{t \in \mathcal{T}(d)}$ share the same market decisions, which implies that these two solutions attain the same objective values in their respective optimization problems. Thus, we can always find a feasible solution of (13) that attains the same objective value as any feasible solution of the reduced trader's problem (IT^r). \square

In summary, the results of this section show that, as long as the end-of-day reservoir levels are fixed a day in advance, the infinite-dimensional trader's problem (IT) is equivalent to the tractable linear program (IT^r), whose size scales linearly with the number A of arcs in the reservoir system and the number H of hours per day. This key insight is formalized in the following theorem.

THEOREM 3. *Under Approximation 1, the optimal values of (IT) and (IT^r) are equal.*

Proof This is an immediate consequence of Propositions 6, 7 and 8. \square

The relations between the various optimization problems studied so far in this appendix are illustrated in the upper row of Figure 8, which uses the same conventions as Figure 1. To conclude the appendix, it remains to be shown that the optimal objective value of (CT^r) is upper bounded by that of (IT^r). This is established in the next and final proposition.

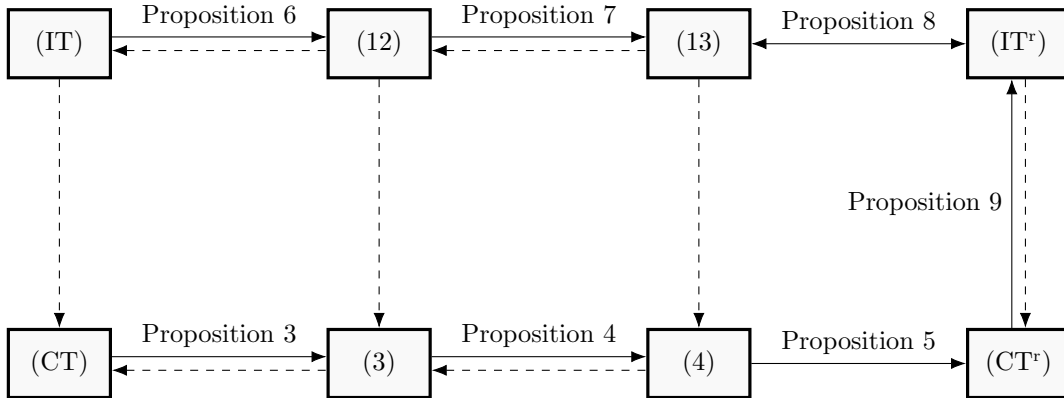


Figure 8 Illustration of the relations between different variants of the trader's problem. Dashed arcs represent trivial relaxations, and solid arcs represent non-trivial implications proved in the referenced propositions.

PROPOSITION 9. *Under Approximation 1, the optimal value of problem (IT^r) is larger than or equal to that of problem (CT^r).*

Proof The claim follows if we can show that for every feasible solution of problem (CT^r) there exists a feasible solution of (IT^r) that attains the same objective function value. To achieve this, we select an arbitrary feasible solution $\{(s_t, \mathbf{u}_t, \mathbf{v}_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ of problem (CT^r) and aim to show that problem (IT^r) admits a feasible solution $\{(s'_t, \mathbf{u}'_t, \mathbf{v}'_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ with the same flow decisions that satisfies $\mathbf{1}^\top s'_t = s_t$, $\mathbf{1}^\top \mathbf{u}'_t = u_t$ and $\mathbf{1}^\top \mathbf{v}'_t = v_t$ for all $t \in \mathcal{T}(d)$. These identities ensure that the two solutions adopt the same objective values in their respective optimization problems.

As the flow decisions $\{(\mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ are preserved, their upper and lower bounds as well as the reservoir level constraints are trivially satisfied. It thus suffices to show that there exist individual market bids $\{(\mathbf{s}'_t, \mathbf{u}'_t, \mathbf{v}'_t)\}_{t \in \mathcal{T}(d)}$ that are consistent with the prescribed collective market bids $\{(s_t, u_t, v_t)\}_{t \in \mathcal{T}(d)}$ and that satisfy all remaining constraints of problem (IT^r), that is, the non-negativity constraints, the energy delivery constraints and the valid inequalities from Proposition 4. Formally, such individual market bids exist if and only if the optimal value of the feasibility problem

$$\begin{aligned}
& \min && 0 \\
& \text{s.t.} && \mathbf{s}'_t \in \mathbb{R}^A, \mathbf{u}'_t \in \mathbb{R}_+^A, \mathbf{v}'_t \in \mathbb{R}_+^A \\
& && \mathbf{1}^\top \mathbf{s}'_t = s_t, \mathbf{1}^\top \mathbf{u}'_t = u_t, \mathbf{1}^\top \mathbf{v}'_t = v_t \\
& && \mathbf{s}'_t + \bar{\rho}_t^u \mathbf{u}'_t = \boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t \\
& && \mathbf{s}'_t - \bar{\rho}_t^v \mathbf{v}'_t \geq -\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t
\end{aligned} \tag{16}$$

vanishes for each hour $t \in \mathcal{T}(d)$, provided that $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ is feasible in (CT^r). Assigning dual variables $\alpha, \beta, \gamma \in \mathbb{R}$ to the consistency constraints for the market bids, $\boldsymbol{\lambda} \in \mathbb{R}^A$ to the energy delivery constraints and $\boldsymbol{\mu} \in \mathbb{R}_+^A$ to the valid inequalities from Proposition 4, the linear program dual to the above feasibility problem can be represented as

$$\begin{aligned}
& \max && \alpha s_t + \beta u_t + \gamma v_t + \boldsymbol{\lambda}^\top (\boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t) - \boldsymbol{\mu}^\top (\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t) \\
& \text{s.t.} && \alpha, \beta, \gamma \in \mathbb{R}, \boldsymbol{\lambda} \in \mathbb{R}^A, \boldsymbol{\mu} \in \mathbb{R}_+^A \\
& && \alpha \mathbf{1} + \boldsymbol{\lambda} + \boldsymbol{\mu} = \mathbf{0} \\
& && \beta \mathbf{1} + \bar{\rho}_t^u \boldsymbol{\lambda} \leq \mathbf{0} \\
& && \gamma \mathbf{1} - \bar{\rho}_t^v \boldsymbol{\mu} \leq \mathbf{0}.
\end{aligned} \tag{17}$$

Strong duality holds because the feasible set of the dual problem contains the origin. In the remainder we will argue that the objective value of any feasible solution $(\alpha, \beta, \gamma, \boldsymbol{\lambda}, \boldsymbol{\mu})$ of the dual linear program (17) is bounded above by 0. To this end, we define $\underline{\mu} = \min_{a \in \mathcal{A}} \mu_a \geq 0$ and note that combining the first two constraints in (17) yields $(\beta - \bar{\rho}_t^u \alpha) \mathbf{1} - \bar{\rho}_t^u \boldsymbol{\mu} \leq \mathbf{0}$, which in turn implies that $\beta - \bar{\rho}_t^u \alpha - \bar{\rho}_t^u \underline{\mu} \leq 0$. We thus find that the objective value in (17) satisfies

$$\begin{aligned}
& \alpha s_t + \beta u_t + \gamma v_t + \boldsymbol{\lambda}^\top (\boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t) - \boldsymbol{\mu}^\top (\boldsymbol{\zeta}_t \circ \bar{\mathbf{p}}_t) \\
& = \alpha s_t + \beta u_t + \gamma v_t + (\boldsymbol{\lambda} + \boldsymbol{\mu})^\top (\boldsymbol{\eta}_t \circ \mathbf{g}_t - \boldsymbol{\zeta}_t \circ \mathbf{p}_t) - \boldsymbol{\mu}^\top (\boldsymbol{\eta}_t \circ \mathbf{g}_t + \boldsymbol{\zeta}_t \circ (\bar{\mathbf{p}}_t - \mathbf{p}_t)) \\
& \leq \alpha s_t + \beta u_t + \gamma v_t - \alpha (\boldsymbol{\eta}_t^\top \mathbf{g}_t - \boldsymbol{\zeta}_t^\top \mathbf{p}_t) - \underline{\mu} (\boldsymbol{\eta}_t^\top \mathbf{g}_t + \boldsymbol{\zeta}_t^\top (\bar{\mathbf{p}}_t - \mathbf{p}_t)) \\
& = \alpha s_t + \beta u_t + \gamma v_t - \alpha (s_t + \bar{\rho}_t^u u_t) - \underline{\mu} (s_t + \bar{\rho}_t^u u_t + \boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t) \\
& = (\beta - \bar{\rho}_t^u \alpha - \bar{\rho}_t^u \underline{\mu}) u_t + \gamma v_t - \underline{\mu} (s_t + \boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t) \\
& \leq (\beta - \bar{\rho}_t^u \alpha - \bar{\rho}_t^u \underline{\mu}) u_t + \underline{\mu} (\bar{\rho}_t^v v_t - s_t - \boldsymbol{\zeta}_t^\top \bar{\mathbf{p}}_t) \leq 0,
\end{aligned}$$

where the first inequality exploits the relations $\boldsymbol{\lambda} + \boldsymbol{\mu} = -\alpha \mathbf{1}$ (by the feasibility of $\alpha, \boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ in (17)) and $\underline{\mu} \mathbf{1} \leq \boldsymbol{\mu}$ (by the construction of $\underline{\mu}$) together with the non-negativity of \mathbf{g}_t and $\bar{\mathbf{p}}_t - \mathbf{p}_t$ (by the feasibility of \mathbf{g}_t and \mathbf{p}_t in (CT^r)). The second equality follows from the energy delivery constraints in (CT^r), and the second inequality holds because $\gamma \leq \bar{\rho}_t^v \underline{\mu}$ (by the feasibility of γ in (17) and the construction of $\underline{\mu}$) and $v_t \geq 0$

(by the feasibility of v_t in (CT^r)). The last inequality, finally, follows from our earlier observations that $\beta - \bar{\rho}_t^u \alpha - \bar{\rho}_t^v \underline{\mu} \leq 0$ and $\underline{\mu} \geq 0$ combined with the relations $u_t \geq 0$ and $\bar{\rho}_t^v v_t - s_t - \zeta_t^\top \bar{\mathbf{p}}_t \leq 0$ (by the feasibility of s_t , u_t and v_t in (CT^r)).

In conclusion, we have demonstrated that the optimal value of the dual feasibility problem (17)—and thus also that of its primal counterpart—must vanish. As our arguments hold for any feasible solution $\{(s_t, u_t, v_t, \mathbf{g}_t, \mathbf{p}_t, \mathbf{z}_t)\}_{t \in \mathcal{T}(d)}$ of (CT^r) and for any $t \in \mathcal{T}(d)$, the claim follows. \square