

The Value of Coordination in Multi-Market Bidding of Grid Energy Storage

Nils Löhndorf

Luxembourg Centre for Logistics and Supply Chain Management, University of Luxembourg, 1511 Luxembourg,
nils.loehndorf@uni.lu

David Wozabal

TUM School of Management, Technische Universität München, 80333 Munich, Germany, david.wozabal@tum.de

We consider the problem of a storage owner who trades in a multi-settlement electricity market comprising an auction-based day-ahead market and a continuous intraday market. We show in a stylized model that a coordinated policy that reserves capacity for the intraday market is optimal and that the gap to a sequential policy increases with intraday price volatility and market liquidity. To assess the value of coordination in a realistic setting, we develop a multi-stage stochastic program for day-ahead bidding and hourly intraday trading along with a corresponding stochastic price model. We show how tight upper bounds can be obtained based on calculating optimal bi-linear penalties for a novel information relaxation scheme. To calculate lower bounds, we propose a scenario tree generation method that lends itself to deriving an implementable policy based on re-optimization. We use these methods to quantify the value of coordination by comparing our policy with a sequential policy that does not coordinate day-ahead and intraday bids. In a case study, we find that coordinated bidding is most valuable for flexible storage assets with high price impact, like pumped-hydro storage. For small assets with low price impact, like battery storage, participation in the day-ahead auction is less important and intraday trading appears to be sufficient. For less flexible assets, like large hydro reservoirs without pumps, intraday trading is hardly profitable as most profit is made in the day-ahead market. A comparison of lower and upper bounds demonstrates that our policy is near-optimal for all considered assets.

Key words: multi-stage stochastic programming, energy storage, electricity price model, limit order book, scenario tree generation, information relaxation

History: published as preprint on November, 30, 2021

1. Introduction

The growing share of renewable power generation along with increasing demand for electric mobility changes electricity systems around the world. Challenges include unforeseen and costly imbalances created by fluctuations in renewable power supply that threaten grid stability as well as surges in demand from electric vehicles that lead to load curtailment and decrease transformer lifetime. Grid energy storage has the potential to address these

challenges by effectively buffering supply and demand and thereby generating significant welfare gains (Sioshansi et al. 2009).

In spite of its benefits and plummeting battery prices, grid energy storage remains scarce (Cole and Frazier 2019, Ziegler et al. 2019). Whether firms view investments in storage as profitable, is largely driven by their expectation of operating cash flows which are difficult to assess and depend on how storage complements an existing portfolio as well as on its ability to improve a company’s trading strategy.

Owners of storage make profit from exploiting inter-temporal price differences which arise due to fluctuating electricity demand as well as varying renewable power generation. Storage assets therefore generate most revenue in short-term markets, where price spreads and thus the value of flexibility is highest.

Power traders that manage electricity storage face a complex decision problem when trading. They can buy or sell in a day-ahead auction for individual hours, blocks of hours, or even quarter-hours of the following day; they can sell spinning and standing reserve capacity in an ancillary service auction; and they can trade continuously in intraday or real-time markets almost until delivery.

In particular, intraday trading has gained traction among European power companies, as the steadily increasing renewable penetration increases the importance of short-term portfolio adjustments due to updated production forecasts. Consequently, the number of commercial software packages to support automated intraday trading is growing steadily (e.g., Likron, Powel Intraday Trading, SOPTIM iTrade, Trayport autoTRADER, etc.).

In this article, we consider the problem of short-term trading in a multi-settlement electricity market with a day-ahead auction and a continuous intraday market. As the number of tradable products in both markets is high and prices evolve randomly, the trader is faced with a multi-dimensional, multi-stage decision problem under uncertainty.

We note that intraday trading is distinct from trading on (real-time) balancing markets that are prevalent in the U.S., as products for individual blocks, hours, and quarter-hours can be traded continuously. This not only creates opportunities for repeated reoptimization, but also leads to a substantial increase in dimensionality of the underlying optimization problem.

To manage the ensuing complexity of the problem, power traders typically follow a sequential strategy, which treats day-ahead bidding and intraday trading as separate decision problems. We address the question of what can be gained from a coordinated trading

strategy, where day-ahead bids anticipate intraday opportunities. The question whether and to what extent a coordinated strategy is better than sequential bidding is of practical importance for power traders who want to maximize the profitability of their storage assets.

Most existing models either focus exclusively on day-ahead bidding or on intraday trading. In Fleten and Kristoffersen (2008), a two-stage stochastic programming model is proposed to find optimal price-dependent bidding curves for a hydropower producer under day-ahead price uncertainty ignoring the intraday market. An extension of the model that accounts for intraday prices is proposed in Faria and Fleten (2011), but the model formulation prohibits allocating capacity to the intraday market at the time of day-ahead bidding. This assumption is dropped in Löhndorf et al. (2013) where short-term bidding is integrated into medium-term planning of pumped-hydro storage, but intraday trading takes place only once and prices are modeled by their expected values, ignoring the value of intraday reoptimization.

Dynamic programming approaches that model the intraday bidding process are proposed in Jiang and Powell (2015) and Aïd et al. (2016), although neither accounts for intraday products with different times to maturity; and while Aïd et al. (2016) model the intraday market, they do not consider storage. Bertrand and Papavasiliou (2019) train a threshold policy for a storage unit on German limit order book data using reinforcement learning. The authors report sizable gains for their trading strategy over a greedy strategy.

Boomsma et al. (2014) extend the work of Fleten and Kristoffersen (2008) and compare sequential with coordinated bidding in day-ahead and balancing markets using multi-stage stochastic programming. Based on data from the Scandinavian Nord Pool market, the authors report gains of 25% for a price-taker and of up to 5% for a market where there is price impact. Kongelf et al. (2019) propose a three-stage stochastic program that additionally considers the market for primary reserves and report gains of less than 1%. As this literature focuses on the Scandinavian markets, where intraday trading still plays only a minor role, intraday trading is not considered. Wozabal and Rameseder (2020) consider the problem of coordinated bidding in sequential auctions for a renewable power producer without storage in the Spanish intraday market and report gains of up to 20%.

In this article, we propose a joint model of day-ahead bidding and intraday trading of storage that considers the option to reoptimize storage and the portfolio of hourly products

multiple times. As opposed to Barbry et al. (2019) who evaluate a price-maker storage and its impact on day-ahead prices in the New York electricity market, we model storage as price-taker in the day-ahead market but recognize the price impact of large orders in the intraday market.

We start by examining the added value of coordinated bidding for a stylized version of the problem and find that, although the intraday market is less liquid, an optimal policy reserves some storage capacity for intraday trading. The reason for the attractiveness of the intraday market is two-fold: first, unlike day-ahead prices, intraday prices can be observed before making a decision to buy or sell. Second, variability of intraday prices is generally higher than day-ahead prices, which leads to higher price spreads that can be exploited by storage. We find that it can even be optimal to allocate all capacity to the intraday market if the market is sufficiently liquid or if the storage asset is small relative to the size of the market.

To assess the value of coordination in a real market situation, we formulate the decision problem as a large multi-stage stochastic program. The model formulation considers hourly day-ahead bids as well as dynamic rebalancing of hourly positions in intraday markets with limited market depth. In accordance with empirical data, we propose a martingale model of day-ahead and intraday prices and model price impact as linear functions estimated from historical settlement data.

Calculating optimal day-ahead bids with this model is not trivial, as it entails the solution of a large-scale multi-stage stochastic optimization problem, which is computationally intractable in its general form. To address this problem, we develop new solution methodology that allows us to calculate lower and upper bounds of the optimal objective value.

Our methodological contribution is three-fold: first, we construct scenario trees for the stochastic optimization problem that lead to lower bounds on the true objective value. The proposed method generates scenario trees that reduce approximation error, preserve the martingale property of the stochastic process, and do not grow exponentially in the number of stages, thereby combining three properties that are necessary to obtain good lower bounds for the considered problem. Second, we show how to improve these lower bounds by a reoptimization heuristic which yields an implementable policy for the storage optimization problem that can be used in practice. Third, in order to be able to assess the quality of our policy, we propose an information relaxation scheme that extends results

in Brown et al. (2010) by computing optimal bi-linear penalties and thereby tight dual upper bounds for the coordinated bidding problem that significantly outperform perfect information bounds. As opposed to Desai et al. (2012), who compute optimal penalties for optimal stopping problems, we use a dualization approach to reformulate the relaxed problem as a single convex stochastic optimization problem with equality constraints. By exploiting specific properties of our problem, we are able to circumvent the nested sampling approach of Desai et al. (2012) and can effectively solve the problem by conventional sample average approximation.

Based on these results, we calculate the value of coordination for three common types of grid energy storage in an out-of-sample case study: a large-scale pumped-hydro storage, a medium-sized hydropower plant with a large reservoir and natural inflow, and a small battery storage. The proposed reoptimization heuristic yields profits that are up to 29.1% higher than those obtained with a sequential trading strategy and that exhibit optimality gaps between 0.0% and 7.2%.

In line with the findings from our stylized model, we find that a bang-bang policy that alternates between full charge and full discharge is not optimal for day-ahead trading, which contrasts results from the literature, e.g., Densing (2013) who studies storage as price-taker. Instead, optimal day-ahead bids minimize the price impact from intraday rebalancing by leaving slack capacity for intraday adjustments and by anticipating the option to reoptimize the portfolio of hourly products.

We identify a sizable value of coordination for large and flexible pumped-hydro storage which has significant price impact in the intraday market. Storage owners with such large and flexible assets therefore benefit most from coordinating day-ahead and intraday bidding decisions.

As predicted by our stylized model, our results further suggest that the small battery storage benefits less from coordination, since the smaller price impact enables batteries to more efficiently exploit intraday price spreads. This finding informs investors in battery storage who should pay more attention to possible cash flows from intraday trading as opposed to merely evaluating storage based on day-ahead price spreads.

We observe the opposite for the less flexible hydropower plant with the large reservoir. The plant hardly benefits from intraday trading at all and makes most of its profits in the day-ahead market. This informs power companies which operate large cascaded reservoirs,

as they are often found in mountainous regions, that little can be gained from including information on intraday trading into their long-term planning models.

Notation : Throughout the paper, we work on a general probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We denote by T the number of decision stages and define a filtration $\{\emptyset, \Omega\} = \mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_T \subseteq \mathcal{F}$ such that all random quantities ξ_t that realize in period t are measurable with respect to \mathcal{F}_t , denoted by $\xi_t \triangleleft \mathcal{F}_t$. We write $\xi^t = (\xi_0, \dots, \xi_t)$ for the *history* of decisions and random variables. Furthermore, for any $N \in \mathbb{N}$, we denote the set $\{1, \dots, N\}$ by $[N]$, the indicator function of set A by $\mathbb{1}_A$, and the inner product in \mathbb{R}^n by $\langle \cdot, \cdot \rangle$.

2. The Value of Coordinated Bidding

In this section, we set up a stylized model that identifies the trade-off faced by a storage owner who participates in the day-ahead and the intraday market. In particular, we show that an optimal decision has to balance the advantages of higher liquidity in the day-ahead market with the gain from having better price information and a greater variability of prices in the intraday market. The model informs the discussion in later parts of the paper, where the effects discussed in this section are quantified in a realistic setting for different types of assets.

We assume that the storage owner is risk-neutral and operates an empty, perfectly efficient storage with energy and power capacity equal to one over two time periods, $t \in \{1, 2\}$. The storage owner is price-taker in the day-ahead market and price-setter in the intraday market. Day-ahead bids q_t^D for delivery in $t \in \{1, 2\}$ are submitted at $t = 0$ before day-ahead prices become known at the beginning of $t = 1$. In $t \in \{1, 2\}$ the storage owner trades on the intraday market, where she can observe the prices of the respective time period before making a decision. We denote by q_t^I the intraday trades for period t and follow the convention that buying electricity leads to positive q_t^D and q_t^I .

Let $P_1^D \triangleleft \mathcal{F}_1$ be the day-ahead price for delivery in period $t \in \{1, 2\}$ which realizes at time $t = 1$ and let $P_t^I \triangleleft \mathcal{F}_t$ be the intraday price for delivery in period $t \in \{1, 2\}$. Assume that the expected intraday price is equal to the (expected) day-ahead price

$$\mathbb{E}[P_s^I | \mathcal{F}_t] = \mathbb{E}[P_s^D | \mathcal{F}_t], \quad \forall s = 1, 2, \forall t = 0, 1, 2.$$

Furthermore, assume that the volume, q_t^I , traded by the storage on the intraday market in period t has a linear price impact, i.e., for a given reference price P_t^I the storage owner

trades at the price $P_t^I + \beta q_t^I$ with $\beta \in [0, \infty)$. A linear price impact is supported by an empirical analysis of market depth in Section 4. Note that P_t^I is the intraday price that realizes if the storage does not participate in the intraday market.

Based on these assumptions, we can restrict our attention to day-ahead bidding strategies that buy a quantity $q_1^D > 0$ for delivery in period 1 and sell the same quantity in period 2, i.e., $q_2^D = -q_1^D$. We furthermore restrict intraday trading in period 1 to buying, i.e., impose the constraint $q = q_1^I = -q_2^I > 0$, in order not to conflate profits from arbitrage between both markets with the value of coordination. In particular, when trading on the intraday market the strategy does not sell power bought on the day-ahead market but rather uses the storage to capitalize on time spreads between intraday prices for periods 1 and 2. As a result, after the day-ahead auction is cleared, a quantity of $0 \leq \bar{q} = 1 - q_1^D \leq 1$ is available for intraday trading.

Once prices P_1^D , P_2^D , and P_1^I have realized, the value of buying $q = q_1^I \leq \bar{q}$ on the intraday market in period $t = 1$ and selling it in $t = 2$, i.e., setting $q_2^I = -q_1^I$, is given by

$$v(q) = -(P_1^I + \beta q)q + \mathbb{E}[(P_2^I - \beta q)q | \mathcal{F}_1] = q(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I) - 2\beta q^2.$$

Maximization in q yields $q^* = \frac{1}{4\beta}(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I)$ and

$$v(q^*) = \frac{1}{4\beta}(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I)^2 - \frac{2\beta}{(4\beta)^2}(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I)^2 = \frac{1}{8\beta}(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I)^2. \quad (1)$$

Note that the price impact β negatively affects the profits of intraday trading in (1). We thus arrive at the following conditional intraday value of storage in period $t = 1$, assuming that the maximal capacity for intraday trading is \bar{q} ,

$$V_1^I(\bar{q}) = \begin{cases} \frac{1}{8\beta}(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I)^2 & \text{if } \mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I < 4\beta\bar{q}, \\ \bar{q}(\mathbb{E}[P_2^I | \mathcal{F}_1] - P_1^I - 2\beta\bar{q}) & \text{otherwise.} \end{cases} \quad (2)$$

Here, the second line covers the case where q^* is larger than \bar{q} and thus infeasible.

We start by proving the following proposition for the case that the day-ahead bids are market orders, i.e., are executed independent of price.

PROPOSITION 1. *If day-ahead bids are independent of price and it is possible that intraday prices in period $t = 1$ are greater than in period $t = 2$ and in particular*

$$\mathbb{P}(P_1^I > P_2^I | P_1^I \geq \mathbb{E}[P_2^I | \mathcal{F}_1]) > 0, \quad (3)$$

then

1. *it is never optimal to allocate all capacity to the day-ahead market;*
2. *there exists a $\bar{\beta} > 0$ such that for a price response $\beta \in [0, \bar{\beta}]$, it is optimal to reserve all capacity for the intraday market and not to participate in the day-ahead auction.*

Proof. We start by noting that a lower bound for the intraday value of storage in $t = 0$ can be obtained by setting $q = \bar{q}$ in (2), which implies that we ignore the effect of the price impact on the optimal decision and always trade the maximum capacity. A risk neutral agent only trades on the intraday market if $P_1^I < \mathbb{E}[P_2^I | \mathcal{F}_1] = P_2^D$ at $t = 1$. From this, we obtain an unconditional lower bound on the value of intraday trading in $t = 0$ by

$$\underline{V}_0^I(\bar{q}) = \mathbb{P}(P_1^I < P_2^D) \bar{q} (\mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] - 2\beta \bar{q}). \quad (4)$$

If $\mathbb{E}[P_1^D] \geq \mathbb{E}[P_2^D]$, there is nothing to show, since the storage owner has no incentive to participate in the day-ahead auction. It is therefore sufficient to focus on the case $\mathbb{E}[P_1^D] < \mathbb{E}[P_2^D]$ where the expected profit from day-ahead trading a quantity $(1 - \bar{q})$ equals

$$V_0^D(1 - \bar{q}) = (1 - \bar{q}) \mathbb{E}[P_2^D - P_1^D]. \quad (5)$$

Note that from our assumptions it follows that $\mathbb{E}[P_1^I] = \mathbb{E}[P_1^D] < \mathbb{E}[P_2^D]$ and consequently $\mathbb{P}(P_1^I < P_2^D) > 0$, which allows us to define

$$\bar{\beta} = -\frac{\mathbb{P}(P_1^I \geq P_2^D)}{4 \cdot \mathbb{P}(P_1^I < P_2^D)} \mathbb{E}[P_2^I - P_1^I | P_1^I \geq P_2^D],$$

such that for price responses $0 < \beta < \bar{\beta}$, it follows

$$\begin{aligned} \frac{\partial}{\partial q} \underline{V}_0^I(q) &= \mathbb{P}(P_1^I < P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] - 4\beta q \mathbb{P}(P_1^I < P_2^D) \\ &> \mathbb{P}(P_1^I < P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] + \mathbb{P}(P_1^I \geq P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I \geq P_2^D] \\ &= \mathbb{E}[P_2^I - P_1^I] = \mathbb{E}[P_2^D - P_1^D] = \frac{\partial}{\partial q} V_0^D(q). \end{aligned}$$

By (3), $\mathbb{E}[P_2^I - P_1^I | P_1^I \geq P_2^D] < 0$ and therefore $\bar{\beta} > 0$. It follows that for $\beta < \bar{\beta}$, the marginal expected profit from trading on the intraday market is higher than the marginal expected profit from trading on the day-ahead market for all $0 \leq q \leq 1$, which means that it is optimal not bid on the day-ahead market and reserve the whole storage capacity for the intraday market. Since this is true for the lower bound \underline{V}_0^I , it also has to hold for V_0^I and the first part of the proposition follows.

To prove the second point, suppose the storage owner allocates $(1 - \bar{q})$ capacity to the day-ahead market and \bar{q} to the intraday market. From (4) and (5), we have a lower bound for the value of storage,

$$\underline{V}(\bar{q}) = (1 - \bar{q})\mathbb{E}[P_2^D - P_1^D] + \mathbb{P}(P_1^I < P_2^D)(\mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D]\bar{q} - 2\beta\bar{q}^2).$$

Maximizing $\underline{V}(\bar{q})$ in \bar{q} yields

$$\underline{q}^* = \frac{1}{4\beta} \left(\mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] - \frac{\mathbb{E}[P_2^D - P_1^D]}{\mathbb{P}(P_1^I < P_2^D)} \right). \quad (6)$$

Note that since we are using a lower bound for the intraday profit, \underline{q}^* underestimates the actual optimal allocation q^* to the intraday market, i.e., $q^* \geq \underline{q}^*$, and therefore overestimates the fraction of storage volume $(1 - \bar{q})$ that would be dedicated to day-ahead trading in optimality.

We now show that $\underline{q}^* > 0$, which establishes that $q^* > 0$, i.e., not all capacity is used in the day-ahead market. If $\underline{q}^* \leq 0$, it is optimal for the storage owner to allocate all capacity to the day-ahead market. Since $\frac{1}{4\beta} > 0$, this requires the second term in (6) to be smaller than or equal to zero, which is equivalent to

$$\begin{aligned} 0 &\geq \mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] \mathbb{P}(P_1^I < P_2^D) - \mathbb{E}[P_2^D - P_1^D] \\ &= \mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] \mathbb{P}(P_1^I < P_2^D) - \mathbb{E}[P_2^I - P_1^I] \\ &= \mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D] \mathbb{P}(P_1^I < P_2^D) - (\mathbb{P}(P_1^I \geq P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I \geq P_2^D] \\ &\quad + \mathbb{P}(P_1^I < P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I < P_2^D]) \\ &= -\mathbb{P}(P_1^I \geq P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I \geq P_2^D] \\ &= -\mathbb{P}(P_1^I \geq P_2^D) \mathbb{E}[P_2^I - P_1^I | P_1^I \geq \mathbb{E}[P_2^I] | \mathcal{F}_1]. \end{aligned}$$

This leads to contradiction, since $\mathbb{E}[P_2^I - P_1^I | P_1^I \geq \mathbb{E}[P_2^I] | \mathcal{F}_1] < 0$ by assumption (3). \square

Note that assumption (3) is always fulfilled if the intraday price for period $t = 2$ is not larger than the intraday price for period $t = 1$ almost surely. This is a rather mild assumption whose violation obviously would be prohibitive for intraday trading.

A coordinated bidding strategy explicitly reserves capacity for the intraday market where price spreads are higher, whereas a sequential strategy allocates all capacity to the day-ahead market and trades on the intraday market by reoptimizing the day-ahead schedule

ex-post. Due to its simple structure the sequential strategy is often adopted by practitioners. Proposition 1 shows that sequential bidding is suboptimal, as it is never optimal to bid the entire capacity on the day-ahead market. Quite the contrary, if the intraday price impact is sufficiently small, it is optimal to abstain from day-ahead trading entirely.

The intuition behind Proposition 1 is that storage owners can observe the intraday price for $t = 1$ before buying, which is not possible in day-ahead trading, where prices for both periods are revealed simultaneously after bids have been placed. It is this information gain that reduces the probability of losses and thus makes intraday trading relatively more attractive as long as the price response is small. Note that this effect also increases in the size of the storage, which has been set to unit capacity for our analysis. It can be easily seen that increasing storage capacity is equivalent to increasing β and therefore results in a relatively less attractive intraday market for larger assets.

Most day-ahead markets accept limit orders in the form of bidding functions, where bid quantity is expressed as a function of price, which reduces the informational advantage of trading on the intraday market. For our next analysis, we will therefore assume that, using bidding functions, this advantage can be made to completely disappear. In particular, we assume that bid quantities can be specified as arbitrary function of realized prices and that the second stage bid can be made conditional on the acceptance of the first stage bid. Note that real world implementations of day-ahead bidding offer much less flexibility and in particular the second assumption leads to an overestimation of profits that can be made on the day-ahead market.

In this setting, it is optimal to trade if the expectation of the price in $t = 2$, $\mathbb{E}[P_2^D|P_1^D]$, is greater the price for $t = 1$, P_1^D , so that the expected profit from trading a quantity q on the day-ahead market is given by

$$\Pi^D(q) = q \cdot \mathbb{E}[(\mathbb{E}[P_2^D|P_1^D] - P_1^D)\mathbf{1}_{B^D}]$$

with $B^D = \{\omega \in \Omega : \mathbb{E}[P_2^D|P_1^D] > P_1^D\}$. Now suppose that intraday prices have a higher dispersion than day-ahead prices, so that extreme intraday prices are more probable than extreme day-ahead prices. Ignoring price response, this may lead to a situation where expected margins on the intraday market are higher than those on the day-ahead market, i.e.,

$$\mathbb{E}[(\mathbb{E}[P_2^I|P_1^I] - P_1^I)\mathbf{1}_{B^I}] > \mathbb{E}[(\mathbb{E}[P_2^D|P_1^D] - P_1^D)\mathbf{1}_{B^D}] \quad (7)$$

with $B^I = \{\omega \in \Omega : \mathbb{E}[P_2^I | P_1^I] > P_1^I\}$. In the next proposition, we establish that even with a positive price response, it is never optimal to allocate all capacity to the day-ahead market as long as (7) is fulfilled.

PROPOSITION 2. 1. *If condition (7) holds and day-ahead bids are arbitrary functions of price, Proposition 1 still holds.*

2. *If (P_1^D, P_2^D) and (P_1^I, P_2^I) are jointly normally distributed with identical means and correlation, and the variances of the intraday prices in period $i = 1, 2$ are larger than the variances of the spot prices, then (7) is fulfilled.*

Proof. Let X and Y be two random variables. For the purpose of this proof we define

$$\Pi(X, Y, q) = q\mathbb{E}[(\mathbb{E}[Y|X] - X)\mathbb{1}_B]$$

with $B = \{\omega \in \Omega : \mathbb{E}[Y|X] > X\}$. Using this notation, we have $\Pi^D(q) = \Pi(P_1^D, P_2^D, q)$. We furthermore denote by $\Pi^I(q)$ the intraday profit of trading a quantity q conditional on period 1 prices.

The condition for buying in the first period on the intraday market becomes

$$P_1^I + q\beta < \mathbb{E}[P_2^I | P_1^I] - q\beta,$$

which for higher values of β leads to less trading and less profits,

$$\Pi^I(q) = q\mathbb{E}[(\mathbb{E}[P_2^I | P_1^I] - P_1^I)\mathbb{1}_{B'}] - 2q^2\beta \geq \underbrace{\Pi(P_1^I, P_2^I, q) - 2q^2\beta}_{=:\underline{\Pi}^I(q)},$$

where $B' = \{\omega \in \Omega : \mathbb{E}[P_2^I | P_1^I] - 2\beta q > P_1^I\}$ is a subset of B^I , taking into account the price response on the intraday market. The right-hand side of the inequality above thus represents a lower bound on intraday profits associated with a suboptimal trading strategy for the first period, which does not take into account the price response when buying in the first period.

Calculating the marginal value of bidding $q = q_1^D = -q_2^D$ on the day-ahead market using $\underline{\Pi}^I$ instead of Π^I , we get

$$\frac{\partial}{\partial q} (\Pi^D(q) + \underline{\Pi}^I(1 - q)) = \Pi(P_1^D, P_2^D, 1) - \Pi(P_1^I, P_2^I, 1) + 4q\beta. \quad (8)$$

Note that by (7), $\Pi(P_1^I, P_2^I, 1) - \Pi(P_1^D, P_2^D, 1) \geq 0$ and therefore for

$$0 \leq \beta < \bar{\beta} = \frac{1}{4}(\Pi(P_1^I, P_2^I, 1) - \Pi(P_1^D, P_2^D, 1)),$$

the marginal value of bidding in the day-ahead market is always smaller than zero, which implies that it is optimal to reserve the entire storage capacity for the intraday market. Similarly, for an arbitrary $\beta > 0$ there always exists a $q > 0$ that is small enough such that (8) is less than zero, which implies that the marginal benefit of trading marginal quantities on the intraday market always exceeds the marginal benefit of trading the remaining capacity on the day-ahead market. This proves the first point of the proposition.

The second point directly follows from the fact that for jointly normally distributed (P_1, P_2) , the profit $\Pi(P_1, P_2, q)$ is increasing in the standard deviations of P_1 and P_2 .

To see this, we denote by σ_1 , σ_2 , and ρ the standard deviations and the correlation of P_1 and P_2 , respectively. For the bivariate normal distribution, the expected price in period 2 is given by $\mathbb{E}[P_2|P_1] = \mathbb{E}[P_2] + \rho \frac{\sigma_2}{\sigma_1} (P_1 - \mathbb{E}[P_1])$, which leads to a profit of

$$q \left(\mathbb{E}[P_2] - \rho \frac{\sigma_2}{\sigma_1} \mathbb{E}[P_1] - P_1 \left(1 - \rho \frac{\sigma_2}{\sigma_1} \right) \right).$$

It is therefore optimal to buy for delivery in the first period if and only if

$$\mathbb{E}[P_2|P_1] > P_1 \Leftrightarrow P_1 < \frac{\mathbb{E}[P_2] - \rho \frac{\sigma_2}{\sigma_1} \mathbb{E}[P_1]}{1 - \rho \frac{\sigma_2}{\sigma_1}} =: v$$

which results in

$$\begin{aligned} \Pi(P_1, P_2, q) &= q \int_{-\infty}^v \left(\mathbb{E}[P_2] - \rho \frac{\sigma_2}{\sigma_1} \mathbb{E}[P_1] - P_1 \left(1 - \rho \frac{\sigma_2}{\sigma_1} \right) \right) \phi_{P_1}(P_1) dP_1 \\ &= q \left(\Phi(\nu) (\mathbb{E}[P_2] - \mathbb{E}[P_1]) + \left(1 - \rho \frac{\sigma_2}{\sigma_1} \right) \phi(\nu) \sigma_1 \right), \end{aligned} \quad (9)$$

where ϕ_{P_1} is the density of P_1 , ϕ and Φ are the density and the distribution function of the standard normal distribution,

$$\nu = \frac{v - \mathbb{E}[P_1]}{\sigma_1} = \frac{\mathbb{E}[P_2] - \mathbb{E}[P_1]}{\sigma_1 - \rho \sigma_2},$$

and the equality in (9) follows from properties of the truncated normal distribution. Noticing that $\frac{\partial}{\partial \sigma_1} \phi(\nu) = -\nu \phi(\nu)$, it is a tedious but straightforward calculation to show that the partial derivative of (9) with respect to σ_1 and σ_2 is strictly positive, i.e., that the profit grows in the variance of P_1 and P_2 . \square

The second part of the proposition illustrates that in a setting with normally distributed prices, coordinated bidding is optimal as long as intraday prices are more volatile than

day-ahead prices, which is a realistic assumption and shows that (7) is a sensible condition. We also note that in practice, even if the prices are not normally distributed, (7) will be easy to verify either by direct calculation or by estimating the involved quantities using Monte-Carlo simulation.

Note that price-response threshold $\bar{\beta}$ (that guarantees that trading takes place only on the intraday market) changes between Proposition 1 and Proposition 2 due to two opposing effects. First, the day-ahead market in Proposition 2 is more attractive than the day-ahead market in Proposition 1, as it allows for the submission of bidding functions. This effect decreases $\bar{\beta}$ in Proposition 2 relative to Proposition 1. Second, condition (7) introduces an additional incentive to trade on the intraday market and $\bar{\beta}$ clearly increases in the difference between the left-hand side and the right-hand side of the equation. Depending on which of these two effects prevails, $\bar{\beta}$ in Proposition 1 can be both larger or smaller than $\bar{\beta}$ in Proposition 2.

Of course, rules for bidding functions on real markets are more restrictive than the setting used in Proposition 2. This implies that in reality bidding functions cannot completely offset the informational advantage of the intraday market which further strengthens the case for coordination.

In summary, the analysis of the stylized model shows that the value of coordination between the intraday and the day-ahead market is driven by the trade-offs between liquidity costs, price volatility, as well as the information gain of trading on the intraday market.

3. A Multi-stage Stochastic Program of Coordinated Bidding

In the previous section, we established that sequential bidding is suboptimal in a stylized setting. In this section, we propose a multi-stage stochastic program to quantify the gains of a coordinated bidding strategy in a much more realistic model of bidding and storage operation. We formulate the model with the European market design in mind, where a combination of day-ahead auction and continuous intraday trading is currently becoming the norm.

3.1. Background on Day-ahead and Intraday Trading

In Europe, market participants can trade electricity as bilateral (over-the-counter) contracts or in power exchanges whose trade volumes have been growing constantly over the past decade.

EPEX SPOT is the largest European power exchange with a total trade volume of 567 terawatt-hours (TWh) in 2018. The exchange hosts day-ahead auctions for hourly as well as quarter-hourly products for several European countries and additionally offers the possibility to trade in continuous intraday markets until shortly before physical delivery.

In the day-ahead market, electricity is traded in a sealed bid, single price daily auction for delivery on the next day. The auction closes around noon of the previous day and market clearing is announced roughly one hour later.

In the intraday market, trading for any given hour opens at 15:00 on the previous day and electricity is traded continuously until five minutes before delivery. As this allows for a high degree of flexibility, market participants use intraday markets to make last minute adjustments in case of changes in electricity production or consumption.

In spite of its increasing importance, the intraday market is not sufficiently liquid to accommodate large volumes. Although the intraday market has been growing steadily, its trade volume of 82 TWh in Germany in 2018 is still small compared to 485 TWh traded in the day-ahead auction in the same period (EPEX 2018). Power companies therefore fix their volumes in the day-ahead auction and then use the intraday market to adjust their portfolio. Such sequential trading strategies are typically based on a combination of deterministic optimization based on price forward curves of day-ahead prices and simple trading rules for intraday order execution.

3.2. Assumptions

As in Section 2 we consider a risk-neutral storage owner with a single asset who is price-taker in the day-ahead market and price-setter in the intraday market. We assume that every day the storage owner submits price-dependent bids to the day-ahead market for each hour of the following day. After day-ahead prices realize the storage owner decides about rebalancing the portfolio of hourly products by trading in the intraday market.

We assume day-ahead and intraday prices to have the same expectations, which means that the price of a particular hour is a martingale. This assumption is supported by our empirical analysis in Section 4 (see Figure 2a), which shows that there are no systematic differences between day-ahead and intraday prices. As discussed in Klaassen (1998), this property is important to avoid bias from speculative trading, which occurs whether or not

trading limits and transaction cost are present in the model. Speculation is also not desired by the regulator who requires day-ahead schedules to be balanced (e.g., BMWi 2006).

Our model of the intraday market follows Bertsimas and Lo (1998) and Almgren and Chriss (2001) who model price dynamics of a continuous market as arithmetic random walk with price impact as a linear function of order quantity. The authors propose models with permanent and temporary price impact. We adopt the argument of Gatheral and Schied (2013) arguing that the price impact is only temporary as prices are forced back to equilibrium values by competitive pressure from one hour to the next. Moreover, we model price impacts as deterministic but varying by time to maturity and product.

We assume that injection and withdrawal limits as well as efficiency of the storage are independent of storage content, i.e., we ignore the effect of differences in water pressure on power capacity and efficiency in pumped-hydro storage systems. Furthermore, we assume that storing electricity is subject to losses.

To simplify matters, we assume a fixed storage level at the beginning of the day that we have to return to at the end of the day. In this way, we isolate the effect of coordinating day-ahead and intraday trading from seasonal profits that are accrued over longer time horizons, as these profits are asset-specific and would distort the comparison of sequential and coordinated decisions. Of course, in practice, short-term planning has to be integrated with long-term planning in some way, e.g., by using *water values* or, more generally, a function that assigns a salvage value to the electricity left in the storage at the end of the day. Such information is especially relevant for seasonal storages and can be derived from long-term models such as Löhndorf et al. (2013).

With initial and final storage levels fixed, the model merely needs to keep track of updates of storage operation and of changes in the portfolio of hourly products. Day-ahead bids must be submitted before 12:00. This decision is modeled to take place at stage $t = 0$. Day-ahead market clearing takes place at 12:45 which is when actual day-ahead volumes and prices become known. This is modeled at stage $t = 1$. Intraday trading of hourly products for the next day starts at 15:00 and ends at 22:30. We assume that decisions on intraday trading are made in hourly time increments at the beginning of each hour. This means that intraday trading begins at stage $t = 2$ (15:00) and ends at stage $T = 34$ (22:00), so that each product can be traded until the hour preceding delivery. The resulting stochastic programming problem thus has 35 stages in total.

3.3. Model Formulation

We model the decision problem of the storage owner as a multi-stage stochastic programming problem of the form

$$\begin{aligned} \max_x \quad & \mathbb{E} \left[\sum_{t=0}^T R_t(x^t, \xi^t) \right] \\ \text{s.t.} \quad & x_t \in \mathcal{X}_t(x^{t-1}), \quad x_t \triangleleft \mathcal{F}_t, \quad \forall t \in \{0, \dots, T\}, \text{ almost surely,} \end{aligned} \quad (10)$$

where $\xi = (\xi_0, \dots, \xi_T)$ is the stochastic process of random prices with $\xi_t : \Omega \rightarrow \mathbb{R}^{d_t}$, $R_t(\cdot, \cdot)$ are the immediate reward functions, x_t are the decision variables in t , and $\mathcal{X}_t(\cdot) \subseteq \mathbb{R}^{n_t}$ are the feasible sets.

In the problem discussed below, the feasible sets \mathcal{X}_t as well as the immediate rewards R_t depend only on x_t and not explicitly on the whole history of the decisions. Similarly, R_t only depends on current and not on past prices, i.e., on ξ_t instead of ξ^t . Combined with the fact that ξ is a Markov process, this allows us to write $R_t(x_t, \xi_t)$ and $\mathcal{X}_t(x_{t-1})$ instead of $R_t(x^t, \xi^t)$ and $\mathcal{X}_t(x^{t-1})$.

3.3.1. Decision variables and feasible set. All $H = 24$ hourly products can be traded for nine hours in the intraday market from 15:00 to 23:00 before products go into delivery one at a time. Denote by $\mathcal{H}_t = \{1 + \max(t - 9, 0), \dots, 24\}$ the set of hours that can be traded in period t and by $y_t \in \mathbb{R}^{|\mathcal{H}_t|}$ the positions held in period t . For better readability, we adopt the convention that y_{th} refers to the position held in period t that is due for delivery in hour h . For example, in period $t = 2$, i.e., at 15:00 one day before delivery, $y_{2,10}$ is the 10th element of vector y_2 which is associated with delivery from 9:00 to 10:00, while in $t = 19$, i.e., at 8:00 on delivery day, $y_{19,10}$ is the 1st element of vector y_{19} . We will use analogous indexing for the other variables in the model. Positive values for y_{th} correspond to a long position while negative values correspond to a short position.

Similarly, define $s_t \in \mathbb{R}^{|\mathcal{H}_t|}$ as projected future storage levels at time t if open positions y_t would be closed without further trading by physically dispatching the storage starting from its current level s_{tt} . The vector of all decision variables is denoted by $x_t = (y_t, s_t)$.

At $t = 0$, the decision-maker submits a piecewise constant bidding function for all hours of the next day. This results in buying or selling a random volume for a random price that both realize once the uniform clearing price of the day-ahead auction P_1 is announced in $t = 1$. We follow the approach of Fleten and Pettersen (2005), Löhndorf et al. (2013) and Wozabal and Rameseder (2020) to model the bidding function. In particular, we

choose a grid of $Z \in \mathbb{N}$ prices g_{h1}, \dots, g_{hZ} for all $h = 1, \dots, 24$ and introduce quantity bids y_{0h0}, \dots, y_{0hZ} as variables in our optimization problem.

For a realized day-ahead price P_{1h} in hour h , the realized volumes y_{1h} are given by

$$y_{1h} = y_{0h0} + \sum_{z=1}^Z H(P_{1h} - g_{hz})(y_{0hz} - y_{0h,z-1}), \quad \forall h \in [24], \quad (11)$$

where H is the heaviside function, i.e., the indicator function of the interval $[0, \infty)$. Since the day-ahead market is the first opportunity to trade on a specific day, the realized volumes y_{1h} equal the initial positions held.

Consequently, in the first period the feasible set can be written as

$$\mathcal{X}_0 = \{x_0 = (y_0, s_0) \mid y_{0h,z-1} \geq y_{0hz}, \forall h \in [24], \forall z \in [Z], s_{0h} = S_0, \forall h \in [24]\},$$

where the constraints on the y_{0hz} enforce monotonicity of the bidding functions and S_0 is the initial storage content. At $t = 1$ the day-ahead price realizes and results in a position y_1 . The feasible set can be written as

$$\mathcal{X}_1(x_0) = \left\{ x_1 = (y_1, s_1) \mid \begin{array}{l} y_{1h} = y_{0h0} + \sum_{z=1}^Z H(P_{1h} - g_{hz})(y_{0hz} - y_{0h,z-1}), \forall h \in [24] \\ s_{1h} = s_{0h} + \sum_{k=1}^h y_{1k}, \forall h \in [24] \end{array} \right\}.$$

At $t = 2$ intraday bidding starts and we enforce that all trading decisions must be backed by a physically feasible schedule of storage operation. We assume the storage to be subject to in-kind losses and to have a round-trip efficiency of η , a maximum injection rate into the storage of \bar{y} , a maximum withdrawal rate from the storage of \underline{y} , and a maximum storage content of \bar{s} . Defining $y_t^- = \min(y_t, 0)$, $y_t^+ = \max(y_t, 0)$, the set of feasible trading and operational decisions in period $t \geq 2$ is given by

$$\mathcal{X}_t(x_{t-1}) = \left\{ (y_t, s_t) \mid \begin{array}{l} s_{th} = s_{t,h-1} - \eta y_{th}^- - y_{th}^+, \forall h \in \mathcal{H}_t, \\ s_{th} \in [0, \bar{s}], \forall h \in \mathcal{H}_t, s_{t,0} = s_{t-1,1}, \\ s_{tH} = S_0, y_{th} \in [-\bar{y}, \underline{y}], \forall h \in \mathcal{H}_t \end{array} \right\}. \quad (12)$$

Note that the forward-looking storage balance constraints in (12) enforce that trades for future periods are backed by an implementable schedule of injections and withdrawals (see Löhndorf and Wozabal 2021). This constraint in particular ensures that the problem has relatively complete recourse, which is required for our greedy benchmark strategy discussed in Section 5.3, which otherwise may not be able to fulfill constraint $s_{TH} = S_0$ in the last stage of the problem.

3.3.2. Immediate reward. Denote $\xi_t = P_t : \Omega \rightarrow \mathbb{R}^{|\mathcal{H}_t|}$ as market prices in period t for all hourly products that are tradable in t , with P_1 as day-ahead prices and P_t , $t = 2, \dots, T$ as intraday prices. Day-ahead prices are uncertain in $t = 0$, so that there is no immediate reward in the first stage and the reward in $t = 1$ is given by

$$R_1(x_1, P_1) = -\langle P_1, y_1 \rangle, \quad (13)$$

which is the sum of day-ahead profits that realize after the result of the day-ahead auction is announced but before the respective hourly product can be traded in the intraday market. Notice that at stage $t = 1$ no decisions are made.

Denote $b_t \in \mathbb{R}^{|\mathcal{H}_t|}$ as the vector of slopes of the price impact in period t for every MWh traded. The immediate reward in periods $t = 2, \dots, T$ is given by the sum of all revenues from intraday trading subject to a linear price impact,

$$R_t(x_t, P_t) = \sum_{h \in \mathcal{H}_t} (P_{th} - b_{th}(y_{t-1,h} - y_{th}))(y_{t-1,h} - y_{th}), \quad t = 2, \dots, T. \quad (14)$$

Due to the linear price impact in the objective, the optimization problem becomes a convex quadratic program.

4. Day-ahead and Intraday Price Model

In this section, we describe a stochastic model of day-ahead and intraday prices for the German EPEX SPOT market that will serve as input to the stochastic optimization problem discussed in the last section.

4.1. Day-ahead Model

The day-ahead model is fitted using historical data of hourly day-ahead prices between 2015-10-01 and 2018-09-30 as published on the EPEX SPOT website.

We use a simple time series regression to remove trend and seasonal components from day-ahead price data. The model captures long-term trends by piecewise-linear segments with breakpoints every seven days, and short-term seasonality by dummy variables for each day of the week, with bank holidays treated as Sundays. A separate regression model is fitted for each hour of the day. The regression model explains the day-ahead price $P_{1,d,h}$ in hour h for the d -th observation in the sample by

$$P_{1,d,h} = \beta_{0h}^{(1)} + \sum_{j=0}^{\lfloor d/7 \rfloor} \beta_{jh}^{(2)} \mathbb{1}_{\{7j+1, \dots, 7(j+1)\}}(d) + \sum_{j=1}^7 \beta_{jh}^{(3)} \mathbb{1}_{\{j\}}(d \bmod 7 + 1) + \varepsilon_{dh}^D, \quad (15)$$




Figure 1 Time series regression of day-ahead electricity prices for a summer day (left) and a winter day (right).

where $\beta_{jh}^{(2)}$ is the local linear trend during week j , $\beta_{jh}^{(3)}$ are the daily coefficients, while ε_{dh}^D are the model errors. Note that, consistent with the notation in Section 3 of the paper, the subscript 1 in $P_{1,d,h}$ indicates that prices are day-ahead prices, i.e., that they realize at stage 1 of the decision process.

We fit the above regression model to prices of every hour, $h = 1, \dots, 24$ and use the empirical distributions on the residuals $\hat{\varepsilon}_h^D$ to construct the scenario trees described in Section 5.

Figure 1 shows a comparison of historical and estimated prices for two selected days in summer and winter 2017. The fitted regression model achieves a mean absolute error of EUR 5.56, which, despite the simplicity of the model, is comparable to the performance of alternative day-ahead price forecasts (Weron 2014).

We note that the goal of the above model is not to compete with hourly price forward curves provided by commercial data providers, but rather to obtain a realistic model of day-ahead forecast errors that helps us to measure the added value of a coordinated bidding strategy.

4.2. Intraday Model

The intraday model is based on M days of historical trade and quotes data of intraday price ticks between 2015-11-30 and 2017-09-30. We group observations by quote hour and delivery hour and use the volume-weighted average price of each group as our intraday

prices. As in Aïd et al. (2016), we model the evolution of the intraday price of hourly products as an arithmetic random walk without drift, i.e.,

$$P_{th} = P_{t-1,h} + \varepsilon_{th}^I, \quad t = 2, \dots, T, \quad h \in \mathcal{H}_t, \quad (16)$$

with realized day-ahead price, P_{1h} , as its initial state. The random walk hypothesis is supported by the augmented Dickey-Fuller test which does not reject the presence of a unit root at $\alpha = 0.1$ in 77% of price time series of hourly products. We therefore assume that errors ε_{th}^I are independent across time periods $t = 2, \dots, T$.

We model the distribution of the innovation ε_{th}^I by a discrete distribution with the centered residuals from (16) as its atoms, which we assign uniform probabilities. In particular, each atom

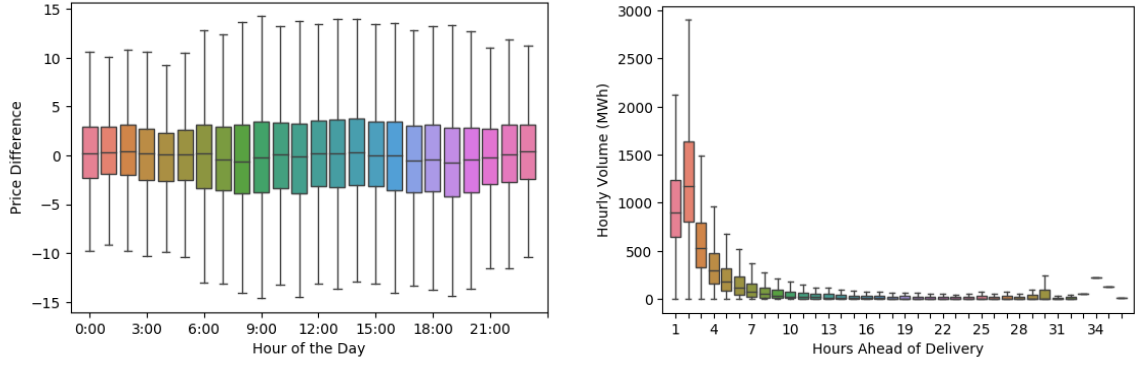
$$\hat{\varepsilon}_{tdh}^I = (P_{t-1,d,h} - P_{tdh}) - M^{-1} \sum_{d'=1}^M (P_{t-1,d',h} - P_{td'h}), \quad d \in [M], \quad t \in \{2, \dots, T\}, \quad h \in \mathcal{H}_t, \quad (17)$$

is assigned probability M^{-1} where, for $t \geq 2$, P_{tdh} are volume-weighted intraday prices for hour h in period t on day d of the decision problem and $P_{1,d,h}$ is the day-ahead price for hour h . Note that the second term in (17) ensures that the distribution has zero mean, which makes the price process a martingale. We additionally model separate distributions for $\hat{\varepsilon}_{tdh}^I$ in summer (quarters 2 and 3) and winter (quarters 1 and 4).

Note that the above formula centers all outcomes of $\hat{\varepsilon}_{tdh}^I$ to ensure that the price process is a martingale. We note that the data shows no significant violation of this assumption. In particular, the martingale assumption is warranted for the difference between day-ahead prices and the volume weighted average of intraday prices one hour ahead of delivery as is illustrated in Figure 2a. Consequently, due to (17), the expected value for the first intraday quotation at 15:00 is assumed to be equal to the realized day-ahead price, which ensures that day-ahead and intraday prices have the same expected values, i.e., the price process is a martingale.

4.3. Intraday Price Impact

To estimate the price impact of trades observed in hour t for a particular product h on day d , we first sort the tick data in t by price in ascending order. We then construct a monotone price impact function by computing the cumulative volumes. Using linear regression, we estimate the relationship between cumulative volume and price for each hour t and product



(a) Deviation of intraday from day-ahead prices (b) Volume of hourly products by time-to-maturity.

Figure 2 Distribution of intraday prices and volumes for hourly products




Figure 3 Settlements of intraday orders for different delivery hours on 2016-12-21 (solid line: function of prices and cumulative volumes in ascending order of price, dotted line: estimated price impact function).

h on day d and construct a sample of slope coefficients b_{tdh} . As we can see in Figure 3, linear functions are reasonable approximations of the relationship between prices and cumulative volumes within one hour. This is in line with the literature on models of financial intraday markets (e.g., Almgren and Chriss 2001, Gatheral and Schied 2013, and references therein).

Figure 2b shows the cumulative hourly trading volume for each hourly product by time-to-maturity. The traded volume sharply drops for hourly products that are more than two hours ahead of delivery. We can also observe this effect in Figure 3 as products that are still four hours away from delivery (orange) exhibit lower trading volumes and steeper price increase than products that are only two hours ahead of delivery (blue). This indicates that the price impact increases in time-to-maturity.




Figure 4 Log-linear regression of observed price impact for products with increasing time-to-delivery during hours 0:00, 4:00, 8:00, 12:00, 16:00, 20:00.

As the distribution of slopes b_{tdh} obtained above is right-skewed, we estimate a model of the linear price impact as a log-linear function of the time-to-maturity of the respective hourly product, i.e.,

$$\log b_{tdh} = \beta_{0t} + \beta_{1t}h + \varepsilon_t^S, \quad j = 1, \dots, T, \quad t = 2, \dots, T, \quad h \in \mathcal{H}_t. \quad (18)$$

Figure 4 shows log-transformed slope coefficients together with the prediction from (18) for a selection of hourly products. The price impact for products that are less than six hours ahead of delivery is negligible, but the regression line has an upward trend, which confirms that the price impact increases in time-to-maturity. For example, after inverse transformation, the product with delivery at 12:00 has an estimated price impact of almost 1.47 EUR/MWh if traded 21 hours ahead of delivery at 15:00.

5. Solution Method

In this section, we discuss how we approximate the stochastic price process from Section 4 by a scenario tree which we use to construct lower bounds of the optimal objective value of the problem introduced in Section 3. Furthermore, we propose a novel upper bound based on information relaxation. We will use these bounds in Section 6 to assess the value of coordination for different types of grid energy storage.

5.1. Scenario Tree Lower Bound

Scenario tree generation is an important topic in stochastic programming and many competing approaches have emerged in the literature during the past decades. Seminal contributions on scenario tree generation aimed at providing lower and upper bounds on the optimal value of the original problem, e.g., Edirisinghe (1996), Frauendorfer (1996), Kuhn (2005), Casey and Sen (2005). Klaassen (1998) and Høyland and Wallace (2001) propose approaches that can be used to construct arbitrage-free trees. More recent approaches use probability metrics to provide bounds on approximation errors, e.g., Pflug (2001), Dupačová et al. (2003), Heitsch and Römisch (2009), Pflug and Pichler (2015).

However, none of the methods are suitable for solving the multi-stage stochastic program in Section 3, as they either construct scenario trees that grow exponentially in the number of time stages, e.g. Frauendorfer (1996) and Høyland and Wallace (2001), or trees that do not preserve the martingale property of the discretized process and thus cannot be used for pricing, e.g. Heitsch and Römisch (2009), or make assumptions that are not fulfilled by our problem, e.g., Casey and Sen (2005) who require that only right-hand sides of constraints are stochastic.

In this section, we propose a simple yet effective algorithm to construct scenario trees that preserve the martingale property of the stochastic process from Section 4 but allows for non-constant branching factor and thereby leads to trees of manageable size that need not grow exponentially in the number of stages by construction. Furthermore, we show that the optimal objectives calculated with the scenario trees are lower bounds of the optimization problems introduced in Section 3.

Since there is a one-to-one correspondence between scenario trees and finite filtrations, we represent the scenario tree by a filtration $\bar{\mathcal{F}} = (\bar{\mathcal{F}}_1, \dots, \bar{\mathcal{F}}_T)$, where $\bar{\mathcal{F}}_t \subseteq \mathcal{F}_t$ and $\bar{\mathcal{F}}_t$ has finitely many elements. Because $\bar{\mathcal{F}}_t$ is finite, it is generated by a partition Ω_t^i of Ω and every Ω_t^i can itself be partitioned by sets $\Omega_{t+1}^j \in \bar{\mathcal{F}}_{t+1}$. Based on the filtration $\bar{\mathcal{F}}$, we define a stochastic process $\bar{\xi} = (\bar{\xi}_1, \dots, \bar{\xi}_T)$ by

$$\bar{\xi}_t = \mathbb{E}[\xi_t | \bar{\mathcal{F}}_t], \quad (19)$$

i.e., $\bar{\xi}_t$ is defined as the process of conditional means on the partitions of Ω implied by $\bar{\mathcal{F}}$.

Using the resulting finite scenario tree approximation $\bar{\xi}$ in the stochastic optimization problem (10), we obtain the following bounding result, which is in some sense dual to

the lower bound for multi-stage stochastic minimization problems with random right-hand sides that Birge and Louveaux (2011) hint to in chapter 10.1.

PROPOSITION 3. *If the functions $\xi^t \mapsto R_t(x^t, \xi^t)$ are concave for all $t \in \{0, \dots, T\}$, then the optimal objective value of (10) with scenario tree $\bar{\xi}$ as stochastic process is a lower bound for the optimal objective value based on the original process ξ .*

Proof. The original stochastic optimization problem in its extensive form can be written as

$$V_0(x_0, \xi_0) := \begin{cases} \max_{x_t} \mathbb{E} \left[\sum_{t=0}^T R_t(x^t, \xi^t) \right] \\ \text{s.t. } x^t \in \mathcal{X}_t(x^{t-1}), x^t \triangleleft \mathcal{F}_t, \quad \forall t \in \{0, \dots, T\}, \text{ almost surely.} \end{cases} \quad (20)$$

Clearly, replacing the measurability constraint $x^t \triangleleft \mathcal{F}_t$ by the stricter condition $x^t \triangleleft \bar{\mathcal{F}}_t$ results in a lower objective value.

Noting that by (19), the concavity of $\xi^t \mapsto R_t(x^t, \xi^t)$, and Jensen's inequality,

$$\mathbb{E} \left[\sum_{t=0}^T R_t(x_t, \xi_t) \right] = \mathbb{E} \left[\sum_{t=0}^T \mathbb{E}[R_t(x^t, \xi^t) | \mathcal{F}_t] \right] \geq \mathbb{E} \left[\sum_{t=0}^T R_t(x^t, \mathbb{E}[\xi^t | \mathcal{F}_t]) \right] = \mathbb{E} \left[\sum_{t=0}^T R_t(x^t, \bar{\xi}^t) \right],$$

we have

$$V_0(x_0, \xi_0) \geq \bar{V}_0(x_0, \bar{\xi}_0) := \begin{cases} \max_{x_t} \mathbb{E} \left[\sum_{t=0}^T R_t(x^t, \bar{\xi}^t) \right] \\ \text{s.t. } x^t \in \mathcal{X}_t(x^{t-1}), x^t \triangleleft \bar{\mathcal{F}}_t, \quad \forall t \in \{0, \dots, T\}, \text{ a.s.} \end{cases} \quad (21)$$

and the result follows. \square

Note that Proposition 3 holds for arbitrary scenario trees, $\bar{\xi}$, that fulfill (19). In order to obtain tight bounds, we propose an algorithm that minimizes the nested distance of Pflug and Pichler (2012). Unlike, for example, algorithms 6 and 7 in Pflug and Pichler (2015) which requires a minimum branching factor, the algorithm proposed below limits the number of nodes at stage t to L_t thereby circumventing exponential growth.

We assume $1 = L_0 \leq L_1 \leq L_2 \leq \dots \leq L_T$ and denote $\mathcal{N}_t = \{1, \dots, L_t\}$ as the indices of nodes at stage t and $\mathcal{S}_{tk} \subseteq \mathcal{N}_{t+1}$ with $|\mathcal{S}_{tk}| \geq 1$ as the set of child nodes in $t+1$ that have node $k \in \mathcal{N}_t$ as parent node, such that $\mathcal{N}_{t+1} = \bigcup_{k \in \mathcal{N}_t} \mathcal{S}_{tk}$.

To minimize the approximation error, the algorithm seeks to minimize the Wasserstein distances between the conditional distributions of the original stochastic process ξ and the scenario tree $\bar{\xi}$. For a comprehensive treatment of the Wasserstein distance and its relation

to optimal scenario tree generation methods we refer to Graf and Luschgy (2000), Heitsch and Römisch (2003), Pflug and Pichler (2014).

For the purpose of this paper, we will focus on the Wasserstein distance of order 2, which can be defined as follows.

DEFINITION 1. Let $\xi : \Omega \rightarrow \mathbb{R}^d$ and $\bar{\xi} : \Omega \rightarrow \mathbb{R}^d$ be two random vectors with image measure Q and \bar{Q} , respectively. The Wasserstein distance of order 2 between ξ and $\bar{\xi}$ is defined as

$$W_2(\xi, \bar{\xi}) = \left\{ \begin{array}{l} \inf_{\pi} \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} \|\xi - \bar{\xi}\|_2^2 \pi(d\xi, d\bar{\xi}) \right)^{\frac{1}{2}} \\ \text{s.t. } \pi(A \times \Omega) = Q(A) \quad \forall A \in \mathcal{A}, \\ \pi(\Omega \times B) = \bar{Q}(B) \quad \forall B \in \bar{\mathcal{A}}, \end{array} \right. \quad (22)$$

where $\|\cdot\|_2$ is the 2-norm in \mathbb{R}^d , the infimum is taken over all probability measures π on $(\mathbb{R}^d \times \mathbb{R}^d, \mathcal{A} \otimes \bar{\mathcal{A}})$, and \mathcal{A} and $\bar{\mathcal{A}}$ are sigma-algebras generated by ξ and $\bar{\xi}$ on \mathbb{R}^d .

Loosely speaking, the Wasserstein distance between two distributions can be defined as the minimal cost of transporting mass from one probability distribution to the other, whereby cost are measured by the distance between the respective outcomes.

To better understand the optimal solution to problem (22) in case one of the distributions is discrete, note that every set of K points $\bar{\xi}_1, \dots, \bar{\xi}_K$ uniquely defines a covering $\Gamma = (\Gamma_1, \dots, \Gamma_K)$ of \mathbb{R}^d by

$$\Gamma_i = \left\{ y \in \mathbb{R}^d : i \in \arg \min_{1 \leq j \leq K} \|y - \bar{\xi}_j\|_2 \right\},$$

which is called a *Voronoi diagram*.

Suppose $\bar{\xi}$ is a discrete probability distribution on $\bar{\xi}_1, \dots, \bar{\xi}_K$ with probabilities $\bar{Q}(\bar{\xi}_i) = Q(\Gamma_i)$ for all $i = 1, \dots, K$. Then, clearly, the best way to transport probability mass from the distribution of ξ to the distribution of $\bar{\xi}$ is to move every point to the closest atom of $\bar{\xi}$, i.e., all points in a Voronoi cell Γ_i to its center $\bar{\xi}_i$, leading to

$$W_2^2(\xi, \bar{\xi}) = \int_{\mathbb{R}^d} \min_{1 \leq k \leq K} \|\bar{\xi}_k - \xi\|_2^2 Q(d\xi).$$

Note that this *transportation plan* is feasible, since we assumed that the probability of $\bar{\xi}_i$ is exactly the probability of the associated Voronoi cell under Q .

Based on this intuition, it can be shown (see Graf and Luschgy 2000, Lemma 3.1) that the problem of finding an optimal discretization of ξ with a fixed number of K atoms, $\bar{\xi}_1, \dots, \bar{\xi}_K$ in terms of the Wasserstein distance is equivalent to solving the K -center problem,

$$\inf_{\bar{\xi}} W_2^2(\bar{\xi}, \xi) = \inf_{\bar{\xi}_1, \dots, \bar{\xi}_K} \int_{\mathbb{R}^d} \min_{1 \leq k \leq K} \|\bar{\xi}_k - \xi\|_2^2 Q(d\xi). \quad (23)$$

Hence, the problem of finding an optimal approximating random variable $\bar{\xi}$ with K atoms can be reduced to the problem of finding K vectors $\bar{\xi}_1, \dots, \bar{\xi}_K$ and then assigning probabilities $Q(\Gamma_i)$ to $\bar{\xi}_i$ to arrive at an approximating random variable $\bar{\xi}$.

We want to use these results to build a tree stage-by-stage such that the conditional distribution for $\bar{\xi}_t$ at a given tree node $\bar{\xi}_{t-1}^l$ is as close to the conditional distribution of ξ_t given that $\xi_{t-1} = \bar{\xi}_{t-1}^l$. Since the K -center problem is NP-hard, we resort to sample average approximation and local optimization techniques. More specifically, given a tree node $\bar{\xi}_{t-1}^l$ with $l \in \mathcal{N}_{t-1}$, we draw a sample, $(\hat{\xi}_t^n)_{n=1}^N$, from ξ_t as described below and then use k-means clustering to obtain a set of K nodes that minimize the distance,

$$\min \left\{ \sum_{n=1}^N \min_{k \in [K]} \|\hat{\xi}_t^n - \bar{\xi}_{tk}\|_2^2 : \bar{\xi}_{tk} \in \mathbb{R}^d, \forall k \in [K] \right\}. \quad (24)$$

The chosen nodes, $\bar{\xi}_{tk}$, $k \in [K]$, imply a partition of the sample,

$$\bar{\Xi}_{tk} = \left\{ \hat{\xi}_t^n | k \in \arg \min_{j \in [K]} \|\hat{\xi}_t^n - \bar{\xi}_{tj}\|_2^2 \right\},$$

where we ignore the possibility that sample points could be at equal distance to two or more nodes and remark that such a situation can be resolved by any tie breaking rule. We define the probability of node k by $q_{tk} = q_{t-1,l} \frac{1}{N} \sum_{n=1}^N \mathbb{1}_{\bar{\Xi}_{tk}}(\hat{\xi}_t^n)$ and note that the resulting discretization fulfills (19).

Algorithm 1 outlines how a scenario tree can be generated stage-by-stage based on these considerations. In line 1, the root node is set to the initial state of the stochastic process ξ_0 . In line 3, for stages $t \in [T]$ and each node $k \in \mathcal{N}_{t-1}$ the algorithm generates a sample $((\hat{\xi}_{tk}^n, \hat{\xi}_{t-1,k}^n))_{n=1}^N$ by first drawing $\hat{\xi}_{t-1,k}^n$ from $\bar{\Xi}_{t-1,k}$ and then sampling $\hat{\xi}_t^k \sim \xi_t | \xi_{t-1} = \hat{\xi}_{t-1,k}^n$. This means that the algorithm draws from partitions associated with nodes from the previous stage to generate a sample of successor outcomes. By setting D_k to ∞ in line 4, the algorithm ensures that all nodes get at least one child node. In lines 6 and 7, the algorithm adds a successor to the node $i \in \mathcal{N}_{t-1}$ with the largest Wasserstein distance between its

```

1 Set  $\bar{\xi}_0^1 \leftarrow \{\xi_0\}$ ,  $\bar{\Xi}_{0,1} \leftarrow \{\xi_0\}$ ,  $q_{0,1} = 1$ 
2 for  $t \leftarrow 1$  to  $T$  do
3   For every  $k \in \mathcal{N}_{t-1}$ , sample  $((\hat{\xi}_{tk}^n, \hat{\xi}_{t-1,k}^n))_{n=1}^N$ 
4   Set  $D_k \leftarrow \infty$ ,  $\mathcal{S}_{t-1,k} \leftarrow \emptyset$  for  $k \in \mathcal{N}_{t-1}$ 
5   for  $j \leftarrow 1$  to  $L_t$  do
6      $i \in \arg \max\{D_k q_{t-1,k}, k \in \mathcal{N}_{t-1}\}$ 
7      $\mathcal{S}_{t-1,i} \leftarrow \mathcal{S}_{t-1,i} \cup \{j\}$ 
8      $\{\bar{\xi}_{tk}\}_{k \in \mathcal{S}_{t-1,i}} \leftarrow \arg \min \left\{ \sum_{n=1}^N \min_{k \in \mathcal{S}_{t-1,i}} \|\hat{\xi}_{ti}^n - \bar{\xi}_{tk}\|_2^2 : \bar{\xi}_{tk} \in \mathbb{R}^d, k \in \mathcal{S}_{t-1,i} \right\}$ 
9      $\bar{\Xi}_{tk} \leftarrow \left\{ \hat{\xi}_{ti}^n : \|\hat{\xi}_{ti}^n - \bar{\xi}_{tk}\|_2^2 \leq \|\hat{\xi}_{ti}^n - \bar{\xi}_{tm}\|_2^2, \forall m \in \mathcal{S}_{t-1,i} \right\}$  for  $k \in \mathcal{S}_{t-1,i}$ 
10     $q_{tk} \leftarrow q_{t-1,i} \sum_{n=1}^N \mathbb{1}_{\bar{\Xi}_{tk}}(\hat{\xi}_{ti}^n)$  for  $k \in \mathcal{S}_{t-1,i}$ 
11     $D_i \leftarrow \sum_{n=1}^N \min_{m \in \mathcal{S}_{t-1,i}} \|\hat{\xi}_{ti}^n - \bar{\xi}_{tm}\|_2^2$ 
12  end
13 end

```

Algorithm 1: Pseudo code of the scenario tree generation algorithm.

child nodes and the samples thereby splitting those partitions that exhibit the largest internal variation. This process is repeated until $|\mathcal{N}_t| = L_t$.

REMARK 1. In the implementation of the algorithm, we correct sample means by shifting the successors of every node by a constant to ensure (19). Therefore, using the tower property, it holds that

$$\mathbb{E}[\bar{\xi}_t | \bar{\xi}_{t-1} = \bar{\xi}_{t-1,j}] = \mathbb{E}[\xi_t | \xi_{t-1} = \bar{\xi}_{t-1,j}], \quad (25)$$

i.e., the process does not distort conditional expectations. It follows that, if ε_t have zero mean and are independent, then by (25)

$$\mathbb{E}[\bar{\xi}_t | \bar{\xi}_{t-1}] = \bar{\xi}_{t-1} + \mathbb{E}[\varepsilon_t] = \bar{\xi}_{t-1}.$$

This property is required to avoid speculation in the energy trading problem outlined in Section 3.

5.2. Reoptimization Heuristic

An improved sampling-based lower bound can be obtained by simulating the process of reoptimizing decisions after each state transition of the stochastic process. Reoptimization

heuristics are common among practitioners and referred to as rolling horizon optimization (Sethi and Sorger 1991), lookahead approximation (Powell 2019), or model predictive control (Schildbach and Morari 2016).

Typically, reoptimization heuristics replace uncertain model parameters with their expected values. We slightly modify this approach by solving the stochastic programming problem using the scenario tree from Section 5.1 to obtain a first-stage solution \bar{x}_1 . For all subsequent stages, we proceed in the usual way by solving a deterministic optimization problem over a receding horizon with prices replaced by their expectations. In particular, the first-stage solution of the stochastic programming problem using the scenario tree is given by

$$\bar{x}_1 \in \arg \max_{x_1 \in \mathcal{X}_1(x_0)} \left\{ \mathbb{E} \left[R_1(x_1, \bar{\xi}_1) + \max_{x_t \in \mathcal{X}_t(x_{t-1}), t > 1} \left\{ \mathbb{E} \left[\sum_{t=2}^T R_t(x_t, \bar{\xi}_t) : x_t \triangleleft \mathcal{F}_t \right] \right\} \right] \right\}. \quad (26)$$

To simulate rewards in later stages, we generate N sample paths, $\hat{\xi}_t^n$, $t \in [T]$, $n \in [N]$, and let $\hat{x}_1^n = \bar{x}_1$, $n \in [N]$. For $t = 2, \dots, T$ and $n \in [N]$ and a given $\hat{\xi}_t^n$ and \hat{x}_{t-1}^n , we solve the following deterministic convex quadratic problem to compute decisions

$$\{\bar{x}_t^n, \bar{x}_{t+1}^n, \dots, \bar{x}_T^n\} \in \arg \max_{x_r \in \mathcal{X}_r(\bar{x}_{r-1}^n)} \left\{ R_t(x_t, \hat{\xi}_t^n) + \sum_{r=t+1}^T R_r(x_r, \mathbb{E}[\xi_r | \hat{\xi}_t^n]) \right\} \quad (27)$$

and set $\hat{x}_t^n = \bar{x}_t^n$. This allows us to estimate a lower bound of the optimal objective value by calculating the sample average of total immediate rewards,

$$N^{-1} \sum_{n=1}^N \sum_{t=1}^T R_t(\hat{x}_t^n, \hat{\xi}_t^n).$$

Note that the above reoptimization heuristics provides an implementable policy for the coordinated bidding problem.

5.3. Sequential Bidding

A sequential bidding strategy can be obtained in a similar way by implementing a greedy policy. Instead of using the entire scenario tree to obtain a first-stage solution, we solve the following two-stage problem,

$$\bar{x}_1 \in \arg \max_{x_1 \in \mathcal{X}_1(x_0)} \left\{ \mathbb{E} [R_1(x_1, \bar{\xi}_1)] \right\}, \quad (28)$$

where the expectation is taken over outcomes at nodes \mathcal{N}_1 . In this way, the option to trade intraday is ignored by the day-ahead bid optimization.

To calculate rewards obtained in later stages, we again simulate the decision process as in the last section, whereby the sequential bidding strategy only maximizes immediate rewards,

$$\hat{x}_t^n = \arg \max_{x_t \in \mathcal{X}_t(\hat{x}_{t-1}^n)} \left\{ R_t(x_t, \hat{\xi}_t^n) \right\}. \quad (29)$$

The sample average of the total immediate rewards now provides an estimate of the profit from sequential bidding.

REMARK 2. Note that a pure intraday bidding strategy can be obtained by skipping the two-stage problem and setting $\bar{x}_1 \equiv 0$. We use this strategy in Section 6.3 to study profits of storage assets that do not participate in the day-ahead auction.

5.4. Information Relaxation Upper Bound

To complement the above lower bounds, we propose an upper bound for the optimization problem based on information relaxation in the spirit of Brown et al. (2010). In particular, we solve the problem

$$\begin{aligned} \bar{V}_1(x_0, \xi_0) &:= \mathbb{E} \left[\sum_{t=1}^T \max_{x_t \in \mathcal{X}_t(x^{t-1})} R_t(x_t, \xi_t) - (w_t(x^t, \omega) - \mathbb{E}[w_t(x^t, \omega) | \mathcal{F}_t]) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{E} \left[\max_{x_t \in \mathcal{X}_t(x^{t-1})} R_t(x_t, \xi_t) - (w_t(x^t, \omega) - \mathbb{E}[w_t(x^t, \omega) | \mathcal{F}_t]) \middle| \mathcal{F}_t \right] \right], \end{aligned} \quad (30)$$

for some functions $w_t : \prod_{i=1}^t \mathbb{R}^{n_i} \times \Omega \rightarrow \mathbb{R}$. Since the maximization operator is inside the expectation, the decisions in (30) can be made knowing all outcomes of ξ until the end of the planning horizon, i.e., $x_t \triangleleft \mathcal{F}_T$. Note that, in the notation of Brown et al. (2010), $\mathcal{G}_t = \mathcal{F}_T$, i.e., we derive our bounds based on the perfect information filtration as a relaxation of the natural filtration of the problem.

The last term in (30) involving w_t acts as a penalty for anticipative decisions: In particular, $w_t \equiv 0$, i.e., no penalty, results in the clairvoyant solution. On the contrary, for arbitrary w_t , optimizing over non-anticipative policies at stage t yields

$$\mathbb{E} \left[\max_{x_t \in \mathcal{X}_t(x^{t-1}), x_t \triangleleft \mathcal{F}_t} R_t(x_t, \xi_t) + (w_t(x^t, \omega) - \mathbb{E}[w_t(x^t, \omega) | \mathcal{F}_t]) \middle| \mathcal{F}_t \right]$$

$$\begin{aligned}
&= \max_{x_t \in \mathcal{X}_t(x^{t-1}), x_t \triangleleft \mathcal{F}_t} \mathbb{E} \left[R_t(x_t, \xi_t) + (w_t(x^t, \omega) - \mathbb{E}[w_t(x^t, \omega) | \mathcal{F}_t]) \mid \mathcal{F}_t \right] \\
&= \max_{x_t \in \mathcal{X}_t(x^{t-1}), x_t \triangleleft \mathcal{F}_t} \mathbb{E} [R_t(x_t, \xi_t) | \mathcal{F}_t]
\end{aligned} \tag{31}$$

thereby reducing (30) to the original problem (10). This implies that for non-anticipative policies there is no penalty and consequently $\bar{V}_1(x_0, \xi_0) \geq V_1(x_0, \xi_0)$, i.e., (30) is an upper bound of the true objective value (see Brown et al. 2010, Brown and Smith 2014, for a more detailed exposition).

In the literature on information relaxation bounds, the functions w_t are typically chosen as linear approximations of expected value functions of the respective problems (e.g. Brown et al. 2010, Brown and Smith 2014, Nadarajah and Secomandi 2018, Balseiro and Brown 2019). This choice is motivated by the fact that if

$$w_t(x^t, \omega) = V_{t+1}(x^t, \xi^{t+1}) := \max_{x_t \in \mathcal{X}_t(x^{t-1}), x_t \triangleleft \mathcal{F}_t} R_{t+1}(x^{t+1}, \xi^{t+1}) + \mathbb{E} [V_{t+2}(x^{t+2}, \xi^{t+2}) | \mathcal{F}_{t+1}],$$

i.e., the value function of the stochastic optimization problem, the optimal solution of problem (30) is the true optimum (Brown et al. 2010). It follows that, if functions w_t are close to the optimal value functions of the problem, (30) can be expected to yield tight bounds. This approach, however, requires linear approximations of the true value functions of the problem, which either have to be derived from heuristic solution approaches or obtained numerically.

Our approach, although motivated by the fact that the value functions V_{t+1} are optimal penalties, does not require explicit specification of value function approximations. In particular, we notice that, for Markovian problem, V_{t+1} depends on x_t and ξ_{t+1} and restrict our attention to bilinear penalty functions of the form

$$w_t(x_t, \xi_{t+1}) = \langle x_t, A_t \xi_{t+1} \rangle \tag{32}$$

for some matrix A_t of fitting dimension.

Instead of manually fixing w_t , we seek to find the optimal bilinear bounding function by minimizing over matrices A_t , i.e., solving the minimax problem

$$\min_{A_t \in \mathbb{R}^{n_t \times d_t}} \mathbb{E} \left[\sum_{t=1}^T \max_{x_t \in \mathcal{X}_t(x^{t-1})} R_t(x_t, \xi_t) - (\langle x_t, A_t \xi_{t+1} \rangle - \mathbb{E}[\langle x_t, A_t \xi_{t+1} \rangle | \mathcal{F}_t]) \right]. \tag{33}$$

The corresponding bounds are exact if V_{t+1} is a bilinear function, i.e., if $V_{t+1}(x_t, \xi_{t+1}) = \langle x_t, A_t \xi_{t+1} \rangle$ for some choice of A_t . If this is not the case, our approach picks penalty functions

that yield the lowest upper bounds within the class of function which are bi-linear in x_t and ξ_{t+1} . In particular, this approach guarantees that the resulting bounds are tighter than the perfect information relaxation bound, i.e., the clairvoyant solution, as $A_t = 0$ results in $w_t \equiv 0$, which is a feasible choice in (33). Note that this property is not guaranteed if the penalty w_t is set manually.

Our approach is similar to the dual bounds proposed in Desai et al. (2012), who find an optimal linear combination of Martingale differences for optimal stopping problems. Similar to our Proposition 4, the authors show that the bounding problems can be solved as convex optimization problems. However, there are two fundamental differences between Desai et al. (2012) and our approach. First, the fact that set of feasible actions in Desai et al. (2012) is discrete and therefore the optimal policy for a pathwise optimization can be easily found by comparing finitely many decisions. Secondly, the authors propose a nested sampling approach which consists of taking outer samples on which the policy is evaluated and inner samples to evaluate conditional value functions for every decision made at every stage of an outer sample path.

In our problem, we are faced with a continuous action space. We deal with this problem by interpreting the problem (33) as a Lagrangian minimax problem associated to a linear maximization problem. This makes it possible to interpret (33) as a modified primal problem, where the outer maximization over the coefficients of the matrices A_t is implicitly taken care of by the introduction of additional expectation constraints.

PROPOSITION 4. *If the distribution of ξ_t is finitely supported for all $t \in [T]$, (33) is equivalent to the problem*

$$\begin{aligned} & \max_{x_t} \mathbb{E} \left[\sum_{t=1}^T R_t(x_t, \xi_t) \right] \\ & \text{s.t. } x_t \in \mathcal{X}_t(x_{t-1}), x_t \triangleleft \mathcal{F}_T, \quad \forall t \in [T] \\ & \mathbb{E}[x_{ti}(\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t])] = 0, \quad \forall t \in [T], i \in [n_t], j \in [d_t]. \end{aligned} \tag{34}$$

Proof. Note that the penalty term for stage t can be re-written as

$$\begin{aligned} \mathbb{E}[\langle x_t, A_t \xi_{t+1} \rangle - \mathbb{E}[\langle x_t, A_t \xi_{t+1} \rangle | \mathcal{F}_t]] &= \mathbb{E}[\langle x_t, A_t (\xi_{t+1} - \mathbb{E}[\xi_{t+1} | \mathcal{F}_t]) \rangle] \\ &= \sum_{i,j} a_{tij} \mathbb{E}[x_{ti} (\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t])] \end{aligned}$$

where the a_{tij} are the components of the matrix A_t .

It follows by the interchangeability principle (see Shapiro et al. 2009, Theorem 7.92) that (33) is equal to

$$\begin{aligned} \min_{A_t} \max_{x_t} \mathbb{E} \left[\sum_{t=1}^T R_t(x_t, \xi_t) \right] + \sum_{i,j} a_{tij} \mathbb{E} [x_{ti} (\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t])] \\ \text{s.t.} \quad x_t \in \mathcal{X}_t(x_{t-1}), x_t \triangleleft \mathcal{F}_T, \quad \forall t \in [T]. \end{aligned} \quad (35)$$

Note that since the ξ_t are finitely supported, the decisions x_t are vectors in a finite dimensional space and the above problem can be interpreted as the Lagrangian of a linear optimization problem. By the application of the minimax theorem for linear programming, we therefore can exchange the min and max operator to obtain (see for example Boyd and Vandenberghe 2004, Section 5.4)

$$\begin{aligned} \max_{x_t} \min_{A_t} \mathbb{E} \left[\sum_{t=1}^T R_t(x_t, \xi_t) \right] + \sum_{i,j} a_{tij} \mathbb{E} [x_{ti} (\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t])] \\ \text{s.t.} \quad x_t \in \mathcal{X}_t(x_{t-1}), x_t \triangleleft \mathcal{F}_T, \quad \forall t \in [T], \end{aligned} \quad (36)$$

for which the second objective term is the Lagrangian version of the constraints

$$\mathbb{E} [x_{ti} (\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t])] = 0, \quad \forall t \in [T], i \in [n_t], j \in [d_t].$$

As usual in the minimax approach to duality the inner minimization forces the constraints to hold, showing that (36) is equivalent to (34). \square

REMARK 3. Proposition 4 can be generalized to random variables ξ_t with continuous distributions. Our model does not require such generalization, since in our case ξ_t are generated by innovations that follow empirical distributions of the regression residuals from (15) and (16). We therefore delegate the exposition of the more general case to future work.

Similar to Desai et al. (2012), the advantage of solving problem (34) is that there is no need to specify a problem-specific penalty. Instead the optimal solution to problem (33) chooses the optimal bilinear penalty, thereby leading to tighter bounds than problem-specific linear penalties.

Note that the optimization over penalty coefficients, A_t , ties all scenarios together so that (34) takes the form of a stochastic optimization problem with expectation constraints instead of a pathwise optimization problem. For non-anticipative decisions, $x_t \triangleleft \mathcal{F}_t$, the constraints are automatically fulfilled since

$$\begin{aligned} \mathbb{E} [x_{ti} (\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t])] &= \mathbb{E} [\mathbb{E} [x_{ti} (\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t]) | \mathcal{F}_t]] \\ &= \mathbb{E} [x_{ti} \mathbb{E} [\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j} | \mathcal{F}_t] | \mathcal{F}_t]] = 0. \end{aligned}$$

To obtain numerical solutions, we resort to sample average approximation, e.g., Wang and Ahmed (2008). Note that, as opposed to Desai et al. (2012), in our case there is no need for an expensive inner sampling to be able to evaluate the conditional expectation $\mathbb{E}[\xi_{t+1,j}|\mathcal{F}_t]$ in the equality constraints. In particular, since the innovations $\hat{\varepsilon}_t$ are centered and we assume them to be stage-wise independent, we get $\mathbb{E}[\varepsilon_{t+1,j}|\mathcal{F}_t] = 0$. It follows that

$$\mathbb{E}[x_{ti}(\xi_{t+1,j} - \mathbb{E}[\xi_{t+1,j}|\mathcal{F}_t])] = \mathbb{E}[x_{ti}\varepsilon_{t+1,j}],$$

which can be approximated by its sampled equivalent

$$\sum_{n \in [N]} x_{ti}^n \hat{\varepsilon}_{t+1,j}^n = 0.$$

6. Numerical Results

In this section, we compare coordinated with sequential bidding based on a case study for storage assets that are traded in the German EPEX SPOT market. We generate several benchmark instances by combining real price data with common types of grid energy storage. We use empirical distributions for $\hat{\varepsilon}^I$ and $\hat{\varepsilon}^D$ described in Section 4 to simulate price sample paths for June 21st (summer) and December 21st (winter) of 2017. We choose these two dates as they mark extreme points in solar power generation and power consumption which leads to different electricity price patterns throughout the day.

We study three different types of storage: a large 1000MW pumped-hydro storage with 8000MWh energy capacity and initial storage level of 2000MWh and with 75% round-trip efficiency; a small and fast 10MW battery storage that can be emptied and filled in one hour with 95% round-trip efficiency and zero initial storage; and a seasonal storage with a relatively large reservoir and 100MW power capacity that cannot be emptied on a single day with continuous deterministic inflow equivalent to 50MW of power.

Note that we assume that the hydro reservoir does not have a pump. The condition $s_{tH} = S_0$ in the model formulation can therefore only be met by matching inflows with the production of energy. As opposed to the other two storages, the reservoir's profit therefore does not depend on price spreads but rather on releasing water during hours with highest prices.

As in Löhndorf et al. (2013), we construct bidding functions with $Z = 3$ breakpoints, which leads to price-dependent bids with four price segments. In pre-tests, the increase in profits using more segments is not significant.




Figure 5 Time series of historical and simulated hour-ahead intraday prices.

All problem instances are solved using the FICO Xpress Barrier Solver on a 16-core Intel Xeon 6130 workstation (2.1GHz, 256GB).

6.1. Scenario Tree Lower Bound

We pursue a data-driven approach and directly use empirical distributions of day-ahead forecast errors and intraday price shifts to create scenario trees.

We control the size of the trees by fixing the number of terminal nodes, L_T , and setting

$$L_t = \lfloor tL_T/T \rfloor, \quad t = 1, \dots, T-1, \quad (37)$$

so that the number of nodes increases approximately linearly in the number of stages; and we use $N = 500$ as the sample size to create successor partitions in Algorithm 1.

Figure 5 compares historical hour-ahead intraday prices during quarters Q2-Q3 for the summer day with prices obtained from simulated sample paths drawn from the scenario tree with $L_T = 10,000$ terminal nodes. The scenario tree produces price scenarios that are heavy-tailed like the original data, but are somewhat smoother, so that the tree will likely underestimate the extrinsic value of storage due to lower volatility in intraday price shifts. This is in line with the result from Section 5.1.

Table 1 shows lower bounds of the objective values of the bidding problem using scenario trees with a varying number of terminal nodes for the two chosen dates and the different types of storage. We note that the lower bounds increase, i.e., get tighter, in the number of nodes L_T . Optimal objective values for $L_T = 10,000$ and $L_T = 30,000$ are close but

Storage	Season	Number of Terminal Nodes					
		100	300	1000	3000	10000	30000
Pumped-hydro	Summer	37,279	38,836	40,456	41,334	41,965	42,131
	Winter	139,728	138,293	140,094	141,391	142,153	142,308
Battery	Summer	313	317	326	330	334	335
	Winter	430	433	446	459	467	469
Reservoir	Summer	48,518	48,586	48,708	48,755	48,791	48,798
	Winter	68,575	68,576	68,638	68,689	68,717	68,725

Table 1 Effect of tree discretization on value of lower bound

CPU times increase from 90 minutes to 6 hours. Trees with $L_T = 100,000$ exceed the memory limit of the workstation. In the following comparative analysis, we therefore use $L_T = 10,000$ as the default number of terminal nodes.

6.2. Lower and Upper Bounds for Coordinated Bidding

Next, we compare the lower bounds obtained from solving the stochastic program with the scenario tree (SP) and the reoptimization heuristic (RH) described in Section 5.2 with the upper bounds from the information relaxation problem (IR) described in Section 5.4 and one from solving the problem under perfect information (PI). Table 2 shows optimal objective values for the coordinated bidding problem obtained by the three approaches. While the SP lower bound is deterministic, objective values of RH, IR, and PI are estimated based on 5,000 scenarios.

The gap between IR and RH is between 0.0% for the battery and 7.2% for the pumped-hydro storage, which indicates that the proposed re-optimization heuristic strategy is near-optimal. The large gaps between IR and PI, which are in the range of 14.0% to 67.6%, indicate that the proposed information relaxation provides substantially tighter upper bounds than solving the problem under perfect foresight. By contrast, the gap between the objective value of SP and RH is between 3.0% for the reservoir and 35.6% for the pumped-hydro storage, which demonstrates that the scenario tree significantly underestimates objective values of the more flexible storage units.

Storage	Season	Avg Objective Value (Std Err)			
		SP (LB)	RH (LB)	IR (UB)	PI (UB)
Pumped-hydro	Summer	41965	65182 (475.7)	70241 (522.1)	127638 (637.7)
	Winter	142153	171039 (607.8)	175894 (682.5)	281052 (980.2)
Battery	Summer	334	598 (1.8)	597 (2.4)	1792 (3.8)
	Winter	467	873 (2.5)	870 (3.4)	2687 (5.7)
Reservoir	Summer	48791	50327 (53.1)	50482 (54.4)	58714 (63.5)
	Winter	68717	70298 (113.0)	70400 (116.6)	83334 (124.6)

Table 2 Lower bounds (LB) and upper bounds (UB) for coordinated bidding

Storage	Season	Market	Avg Profit (EUR)			Avg Volume (MWh)		
			Intraday	Sequential	Coordinated	Intraday	Sequential	Coordinated
Pumped-hydro	Summer	Day-ahead	0	25,122	-22,800	0	0	0
		Intraday	38,372	25,355	87,983	14,179	10,574	10,834
		Both	38,372	50,476	65,182	14,179	10,574	10,834
	Winter	Day-ahead	0	128,765	81,473	0	0	0
		Intraday	94,829	26,779	89,567	21,479	9,040	9,070
		Both	94,829	155,544	171,039	21,479	9,040	9,070
Battery	Summer	Day-ahead	0	293	111	0	0	0
		Intraday	513	258	487	320	275	154
		Both	513	551	598	320	275	154
	Winter	Day-ahead	0	387	129	0	0	0
		Intraday	783	418	744	334	310	159
		Both	783	806	873	334	310	159
Reservoir	Summer	Day-ahead	0	47,997	46,864	0	0	0
		Intraday	35,681	1,556	3,463	3,064	1,106	1,099
		Both	35,681	49,554	50,327	3,064	1,106	1,099
	Winter	Day-ahead	0	68,045	67,009	0	0	0
		Intraday	49,976	1,551	3,289	2,718	742	675
		Both	49,976	69,595	70,298	2,718	742	675

Table 3 Average profits and volumes of different trading strategies for different storage assets

6.3. Out-of-Sample Profits and Trade Volumes

We compare expected profits and average trade volumes of sequential and coordinated bidding as well as pure intraday bidding (which is based on sequential bidding) by simulating $N = 10,000$ out-of-sample profits for each strategy.

Table 3 reports lower bounds of expected profits and averages of total buy and sell decisions for the three storage assets on the two chosen dates. The numbers shows that, overall, coordinated bidding performs consistently better than sequential bidding, and sequential bidding performs better than intraday bidding alone. The large pumped-hydro

storage and the battery benefit most from a coordinated bidding strategy, whereas gains for the less flexible reservoir are marginal.

Except for the reservoir, coordinated bidding leads to a noticeable decrease in day-ahead trading volume, which is in line with the result from the stylized model. The coordinated strategy thereby reserves capacity for the intraday market and accepts lower day-ahead profits (even losses) which are offset by higher intraday profits.

The pumped-hydro storage with its high operational flexibility and market impact benefits most from coordinated bidding with a profit increase of 10.0-29.1% over sequential bidding. Clearly, for this asset, the pure intraday strategy performs worst with a loss of 41-45% of profits relative to coordinated bidding. The overall trade volumes are also less for this strategy, since the price response on the intraday market limits the potential for profitable trading. Portfolios that contain large and flexible storage assets should therefore coordinate their bidding decisions across the two markets.

The battery storage gains 8.3-8.5% from a coordinated strategy on average relative to sequential bidding but also only loses 10.3-14.2% of its profits if it does not bid on the the day-ahead market and resorts to a greedy intraday strategy. Given the relatively large share of intraday trading with the battery storage, it may be an option for small storage assets to skip the day-ahead auction entirely and focus on coordinated intraday trading instead – a setting that has also been studied in Bertrand and Papavasiliou (2019).

The large hydro reservoir, by contrast, hardly benefits at all from even entering the intraday market as it makes most of its profits day-ahead. It should be noted that day-ahead volumes are the same for coordinated and sequential bidding as both have to sell the same amount of energy equivalent inflows due to the constraints imposed in (12). Focusing on day-ahead price volatility as in Löhndorf et al. (2013) therefore seems sufficient when optimizing such large and inflexible storages.

Another observation is that the value of storage is generally higher on the winter day than on the day in summer, which can be explained by the overall lower average price spread between day and night hours of only EUR 16.10 on June 21st compared to EUR 35.34 on December 21st.

6.4. Analysis of Trading Decisions

We study the difference in day-ahead bids and intraday trading activity under each bidding strategy.

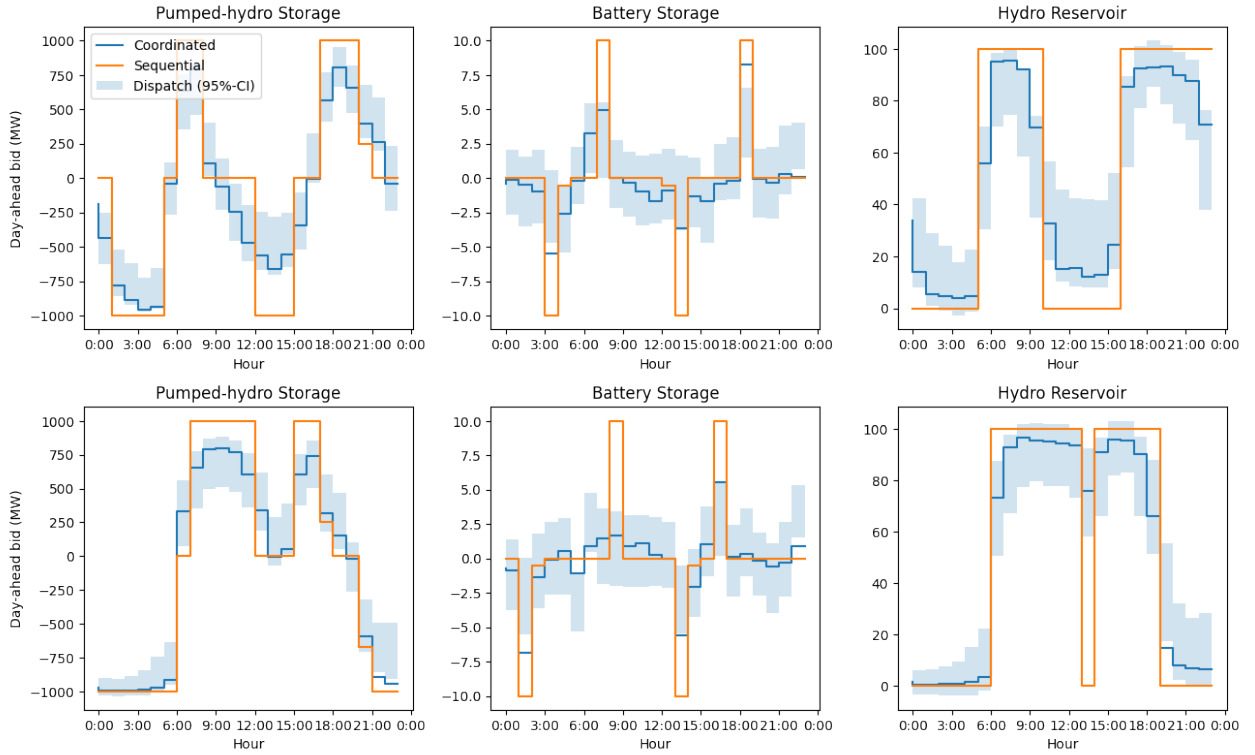


Figure 6 Average day-ahead volumes for each strategy as well as 95% confidence bands of final contract positions (top: summer, bottom: winter).

Figure 6 shows average day-ahead volumes for sequential as well as coordinated bidding for the two chosen days. Unlike sequential bidding, coordinated day-ahead bids are non-trivial and do not fully use the storage’s power capacity, i.e., are not bang-bang solutions. Instead the coordinated strategy places day-ahead bids that match the final intraday position much better, as we can see from their location inside the blue bands in Figure 6 that contain the final positions for 95% of the sampled scenarios.

Figure 7 shows heatmaps of the average order volume in each hour for each product over time, with trading volumes of sequential bidding shown in the top row and coordinated bidding in the bottom row. Coordinated bidding not only shifts more volume into the intraday market but also increases trading volumes as delivery approaches. However, intraday trades are not merely postponed until the last hour but are rather split up taking into account the price impact of individual hours. In contrast, sequential bidding trades earlier on average and therefore is more affected by the price response.

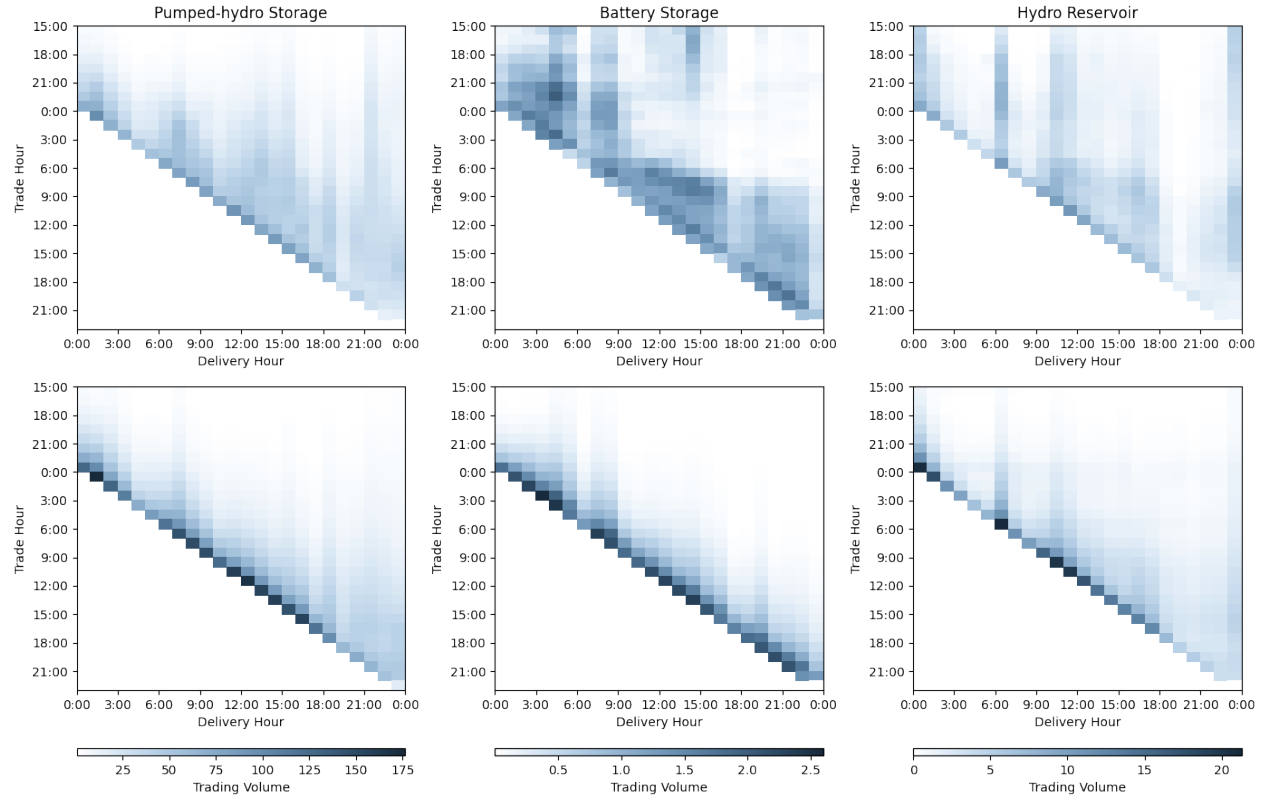


Figure 7 Intraday trading volumes by product and trade hour for different asset types on the summer day (top: sequential bidding, bottom: coordinated bidding).

7. Conclusion

We have modeled the problem of day-ahead and intraday bidding of a storage owner who participates in a multi-settlement market. For a stylized setting, we show that it is optimal to reserve capacity for the intraday market and that a coordinated bidding strategy allows storage owners to make more profit than a sequential strategy.

To measure the added value of coordinated bidding in a real-world setting, we propose a multi-stage stochastic program of day-ahead and intraday bidding. We show how scenario trees can be constructed that allow to compute lower bounds of the optimal objective value and propose a novel upper bound based on information relaxation.

We show in an out-of-sample case study that the added value of using a coordinated bidding strategy is substantial, and that the gain is mostly driven by storage flexibility and its price impact in the intraday market. The value of coordination is not equally large for all asset types. For example, a large and flexible pumped-hydro storage benefits more

from a coordinated bidding strategy than a small battery storage or a less flexible hydro reservoir with natural inflow.

In particular, small and flexible battery storages may opt to skip the day-ahead auction entirely and implement simpler trading strategies that exploit price spreads between intraday products with different times to delivery.

By contrast, little is lost for owners of less flexible hydropower plants with large reservoirs and natural inflow if they do not participate in the intraday market. Quite the contrary, operational cost may even decrease as owners no longer need to maintain and operate costly intraday trading systems.

We also observe that anticipating the intraday market in day-ahead bidding results in complex day-ahead bids that deviate from the bang-bang type solutions often described in the literature.

Our results rest on the assumption that prices for the same hourly products follow a martingale process. By contrast, Ito and Reguant (2016) find empirical support for increasing prices, which creates another strong incentive for market participants to coordinate bidding decisions across time. Whether and in which way our results change if such arbitrage opportunities exist is left to future work.

A natural extension of the optimization model introduced in this paper would be to optimize day-ahead bids and intraday trades of a whole portfolio of production assets including renewables or thermal generators. Such a model could be used by power companies directly to compute day-ahead bids for their entire portfolio. Future work could also investigate formulations where the dynamics of the intraday limit order book is modeled as a continuous-time stochastic process (Cont et al. 2010).

Acknowledgements

The authors would like to thank the Associate Editor and the two anonymous referees for their valuable input and discussion.

References

- Aïd R, Gruet P, Pham H (2016) An optimal trading problem in intraday electricity markets. *Mathematics and Financial Economics* 10(1):49–85.
- Almgren R, Chriss N (2001) Optimal execution of portfolio transactions. *Journal of Risk* 3:5–40.

- Balseiro SR, Brown DB (2019) Approximations to stochastic dynamic programs via information relaxation duality. *Operations Research* 67(2):577–597.
- Barbry A, Anjos M, Delage E, Schell K (2019) Robust self-scheduling of a price-maker energy storage facility in the new york electricity market. *Energy Economics* 78:629 – 646.
- Bertrand G, Papavasiliou A (2019) Adaptive trading in continuous intraday electricity markets for a storage unit. *IEEE Transactions on Power Systems* .
- Bertsimas D, Lo AW (1998) Optimal control of execution costs. *Journal of Financial Markets* 1(1):1–50.
- Birge J, Louveaux F (2011) *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering (Springer New York).
- BMWi (2006) Stromnetzzugangsverordnung (StromNVZ).
<https://www.bmwi.de/Redaktion/DE/Gesetze/Energie/StromNZV.html>.
- Boomsma TK, Juul N, Fleten SE (2014) Bidding in sequential electricity markets: The nordic case. *European Journal of Operational Research* 238(3):797–809.
- Boyd S, Vandenberghe L (2004) *Convex Optimization*. Number Teil 1 (Cambridge University Press), ISBN 9780521833783.
- Brown D, Smith J (2014) Information relaxations, duality, and convex stochastic dynamic programs. *Operations Research* 62(6):1394–1415.
- Brown DB, Smith JE, Sun P (2010) Information relaxations and duality in stochastic dynamic programs. *Operations Research* 58(4):785–801.
- Casey MS, Sen S (2005) The scenario generation algorithm for multistage stochastic linear programming. *Mathematics of Operations Research* 30:615–631.
- Cole W, Frazier A (2019) Cost projections for utility-scale battery storage. URL <https://www.nrel.gov/docs/fy19osti/73222.pdf>.
- Cont R, Stoikov S, Talreja R (2010) A stochastic model for order book dynamics. *Operations Research* 58(3):549–563.
- Densing M (2013) Dispatch planning using newsvendor dual problems and occupation times: application to hydropower. *European Journal of Operational Research* 228(2):321–330.
- Desai V, Farias V, Moallemi C (2012) Pathwise optimization for optimal stopping problems. *Management Science* 58(12):2292–2308.
- Dupačová J, Gröwe-Kuska N, Römisch W (2003) Scenario reduction in stochastic programming. *Mathematical Programming* 95(3):493–511.
- Edirisinghe NCP (1996) Bound-based approximations in multistage stochastic programming: Nonanticipativity aggregation. *Annals of Operations Research* 85:103–127.

- EPEX (2018) Annual report 2018. URL https://www.epexspot.com/sites/default/files/download_center_files/EPEX_RA_WEB%202-compressed.pdf.
- Faria E, Fleten SE (2011) Day-ahead market bidding for a nordic hydropower producer: taking the Elbas market into account. *Computational Management Science* 8(1-2):75–101.
- Fleten SE, Kristoffersen TK (2008) Short-term hydropower production planning by stochastic programming. *Computers & Operations Research* 35(8):2656–2671.
- Fleten SE, Pettersen E (2005) Constructing bidding curves for a price-taking retailer in the norwegian electricity market. *IEEE Transactions on Power Systems* 20(2):701–708.
- Frauentorfer K (1996) Barycentric scenario trees in convex multistage stochastic programming. *Mathematical Programming* 75:277–293.
- Gatheral J, Schied A (2013) Dynamical models of market impact and algorithms for order execution. Jean-Pierre Fouque JAL, ed., *Handbook on Systemic Risk*, 579–599 (Cambridge University Press).
- Graf S, Luschgy H (2000) *Foundations of Quantization for Probability Distributions*. Lecture Notes in Mathematics (Springer Berlin Heidelberg).
- Heitsch H, Römisch W (2003) Scenario reduction algorithms in stochastic programming. *Computational Optimization and Applications* 24(2-3):187–206.
- Heitsch H, Römisch W (2009) Scenario tree modeling for multistage stochastic programs. *Mathematical Programming* 118:371–406.
- Høyland K, Wallace SW (2001) Generating Scenario Decision Trees for Multistage Problems. *Management Science* 47(2):295–307.
- Ito K, Reguant M (2016) Sequential markets, market power, and arbitrage. *American Economic Review* 106(7):1921–57.
- Jiang DR, Powell WB (2015) Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS Journal on Computing* 27(3):525–543.
- Klaassen P (1998) Financial asset-pricing theory and stochastic programming models for asset-liability management: a synthesis. *Management Science* 44(1):31–48.
- Kongelf H, Overrein K, Klæboe G, Fleten SE (2019) Portfolio size’s effects on gains from coordinated bidding in electricity markets. *Energy Systems* 10(3):567–591.
- Kuhn D (2005) *Generalized Bounds for Convex Multistage Stochastic Programs* (Springer-Verlag Berlin Heidelberg).
- Löhndorf N, Wozabal D, Minner S (2013) Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research* 61(4):810–823.
- Löhndorf N, Wozabal D (2021) Gas storage valuation in incomplete markets. *European Journal of Operational Research* 288(1):318 – 330.

- Nadarajah S, Secomandi N (2018) Merchant energy trading in a network. *Operations Research* 66(5):1304–1320.
- Pflug G (2001) Scenario tree generation for multiperiod financial optimization by optimal discretization. *Mathematical Programming, Series B* 89(2):251–271.
- Pflug GC, Pichler A (2012) A distance for multistage stochastic optimization models. *SIAM Journal on Optimization* 22(1):1–23.
- Pflug GC, Pichler A (2014) *Multistage Stochastic Optimization* (Springer Series in Operations Research and Financial Engineering).
- Pflug GC, Pichler A (2015) Dynamic generation of scenario trees. *Computational Optimization and Applications* 62(3):641–668.
- Powell WB (2019) A unified framework for stochastic optimization. *European Journal of Operational Research* 275(3):795–821.
- Schildbach G, Morari M (2016) Scenario-based model predictive control for multi-echelon supply chain management. *European Journal of Operational Research* 252(2):540–549.
- Sethi S, Sorger G (1991) A theory of rolling horizon decision making. *Annals of Operations Research* 29(1):387–415.
- Shapiro A, Dentcheva D, Ruszczyński A (2009) *Lectures on Stochastic Programming: Modeling and Theory*. MOS-SIAM series on optimization (Society for Industrial and Applied Mathematics).
- Sioshansi R, Denholm P, Jenkin T, Weiss J (2009) Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics* 31(2):269 – 277.
- Wang W, Ahmed S (2008) Sample average approximation of expected value constrained stochastic programs. *Operations Research Letters* 36(5):515 – 519.
- Weron R (2014) Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* 30(4):1030–1081.
- Wozabal D, Rameseder G (2020) A stochastic optimization approach for optimal bidding of a virtual power plant on the spanish spot market for electricity. *European Journal of Operational Research* 280(2):639–655.
- Ziegler MS, Mueller JM, Pereira GD, Song J, Ferrara M, Chiang YM, Trancik JE (2019) Storage requirements and costs of shaping renewable energy toward grid decarbonization. *Joule* 3(9):2134 – 2153.