Secondary Publication



Schaurecker, David; Wozabal, David; Löhndorf, Nils; u. a.

High-frequency intraday trading for battery storages

Date of secondary publication: 30.10.2026

Version of Record (Published Version), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-117799x

Primary publication

Schaurecker, David; Wozabal, David; Löhndorf, Nils; u. a. (2026): High-frequency intraday trading for battery storages, in: Applied energy, Amsterdam [u.a.]: Elsevier, vol. 425, no. 128550, pp. 1–15, doi: 10.1016/j.apenergy.2026.128550

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:
https://creativecommons.org/licenses/by/4.0/legalcode

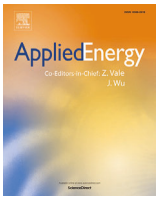
Applied Energy 425 (2026) 128550



Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apen



High-frequency intraday trading for battery storages

David Schaurecker[a,*], David Wozabal[b], Nils Löhndorf[c], Thorsten Staake[d]

[a] ETH Zurich, Rämistrasse 101, 8092 Zurich, Switzerland
[b] Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands
[c] University of Luxembourg, 6 Rue Richard Coudenhove-Kalergi, L-1359, Luxembourg, Luxembourg
[d] University of Bamberg, An der Weberei 5, 96047 Bamberg, Germany

HIGHLIGHTS

- We propose a novel continuous intraday high-frequency trading strategy for BESS.
- Our DP approximation is 100-1000x faster than MILP, maintaining profit.
- High-frequency trading boosts profits: +58% over hourly re-optimization.
- Parametric RI extension increases yearly revenues by 8.4% out of sample.
- We publish the Python package “bitePy” to run high-frequency simulations.

ARTICLE INFO

Keywords:
Electricity trading
Asset-backed trading
Short-term power markets
Dynamic programming
Limit order book

ABSTRACT

Maximizing revenue for grid-scale battery energy storage systems in continuous intraday electricity markets requires strategies that are able to seize trading opportunities as soon as new information arrives. This paper introduces and evaluates a computationally efficient, high-frequency implementation of the rolling intrinsic trading strategy for battery energy storage systems on the intraday market for power. By combining the established rolling intrinsic logic with a full limit order book representation and a fast dynamic programming approximation, our method explicitly considers the continuously updated list of buy and sell offers, market rules, a linear approximation of degradation, and other technical parameters at a millisecond resolution. The standard rolling intrinsic strategy is adapted for continuous intraday electricity markets and solved using a dynamic programming approximation that is two to three orders of magnitude faster than an exact mixed-integer linear programming solution. A detailed backtest over a full year of German order book data demonstrates that the proposed dynamic programming formulation does not reduce trading profits and enables the policy to react to every relevant order book update, enabling realistic rapid backtesting. Our results show the significant revenue potential of high-frequency trading: our policy earns 58 % more than when re-optimizing only once every hour and 14 % more than when re-optimizing once per minute, highlighting that profits critically depend on trading speed. Furthermore, we leverage the speed of our algorithm to train a parametric extension of the rolling intrinsic, increasing yearly revenue by 8.4% out of sample.

1. Introduction

Grid-scale battery energy storage systems (BESS) have emerged as a promising solution to address short-term variability in renewable electricity production. Batteries store excess energy generated during periods of high production and release it when generation is low, providing much-needed flexibility and stabilizing the grid [29,36]. In addition, they help manage peak demand, minimize reliance on fossil fuel peaker plants, and boost overall energy system efficiency. The rapid increase in grid-scale storage installations underscores this potential [19,33], with the projected adoption of grid-scale BESS expected to increase significantly by 2030.

In liberalized electricity markets, the evaluation of viable business models for BESS is crucial to this transformation. The economic feasibility and revenue streams of these systems depend on their ability to participate in various short-term energy markets, engaging in energy arbitrage, frequency regulation, and demand response. Understanding

* Corresponding author.
Email address: schaureckerd@ethz.ch (D. Schaurecker).

https://doi.org/10.1016/j.apenergy.2026.128550
Received 11 March 2026; Received in revised form 16 July 2026; Accepted 26 July 2026
Available online 31 July 2026
0306-2619/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Nomenclature

Sets and Indices

$\mathcal{T}$	Set of delivery time periods / futures contracts
t	Index for time period / delivery contract
t^*	Current time / decision time
t_0	Time when the earliest tradable contract goes into delivery
$\mathcal{O}_t$	Set of all limit orders in the order book for contract t
$\mathcal{O}_t^+, \mathcal{O}_t^-$	Set of currently active asks (+) and bids (–) for contract t
i, j	Indices for individual limit orders
G	Discretized finite grid of possible storage states

Parameters and Constants

$\bar{s}$	Maximum storage energy capacity [MWh]
s_0	Initial storage level [MWh]
$\bar{f}, \underline{f}$	Physical bounds for maximum injection and withdrawal [MWh]
f_t^0	Initial futures positions [MWh]
P_i	Limit price of order i [€/MWh]
Q_i	Available quantity of order i [MWh]
η^+, η^-	Efficiency factors for charging (+) and discharging (–)
ν_{trade}	Volume-based trading cost [€/MWh]
ν_{deg}	Linear degradation cost [€/MWh]
ν	Combined variable cost ($\nu_{\text{trade}} + \nu_{\text{deg}}$) [€/MWh]
u	Minimum tradable unit/quantity [MWh]
m	Number of grid points in the discretized storage state grid G
κ	Action discretization step parameter
ϕ	Penalty parameter for trading products with large bid-ask spreads
$\delta_{t^*,t}$	Bid-ask spread of product t at time t^* [€/MWh]

Variables and Functions

s_t	Storage level at time t [MWh]
f_t	Futures position for contract t [MWh]
f_t^+, f_t^-	Total to be bought (+) and sold (–) volumes for contract t [MWh]
i_t, w_t	Total bought (+) and sold (–) volumes for contract t including previous positions [MWh]
q_t, q_i	Traded quantity of a contract or specific order i [MWh]
k_i, k_t	Integer multipliers for minimum tradable units
α_t	Binary decision variable preventing simultaneous charge/discharge
$V_t(s_{t-1})$	Dynamic programming value function at stage t
$\tilde{V}_t(s_t)$	Linearly interpolated approximation of the value function
π_t	Function encoding current order book cost/revenue information
$S(s_{t-1}, f_t)$	State transition function yielding the subsequent storage level

the potential of these business models helps investors and legislators make the right decisions for a sustainable energy future and increases the adoption of storage technologies. In particular, informed trading decisions for intraday (ID) trading become increasingly relevant, with ID trading volumes growing each year [21] and market participants shifting their trading towards short-term markets.

Finding good trading strategies requires extensive backtesting over long time horizons. However, backtesting the continuous ID market is computationally expensive. Its millisecond frequency generates millions of decision points, making long-term simulations challenging.

This paper introduces a highly realistic, high-frequency implementation of the rolling intrinsic trading strategy for battery energy storage systems. By combining the rolling intrinsic logic with a full limit order book representation and a fast dynamic programming approximation, this approach integrates the full granular evolution of the market with exchange regulations and physical storage constraints. As we argue below, literature on the subject is limited, as most authors adopt a rather stylized view of the intraday market, avoiding most of its complexity and thereby yielding suboptimal strategies and a biased picture of revenue potentials.

Most relevant related literature on coordinated bidding in electricity spot markets focuses on the day-ahead market, such as Fleten and Kristoffersen [12], or treats the intraday market (IDM) as a single trading decision stage [11,22,27]. Applications range from pure trading strategies to asset-focused approaches, such as for hydro plant or battery system operation. Some publications consider a small number of repeated trading decisions in intraday markets for storages [28] and other applications [1,23,37].

Interest in short-term trading strategies for grid-scale BESS has surged recently, a topic explored by, e.g., Sage et al. [32], Cornejo et al. [8] and Zhang et al. [39].

Literature that exclusively focuses on (automated) trading strategies for the intraday market, while capturing its full complexity, is scarce. Some recent work presents approaches for single-, or multi-market trading strategies for storage. Jiang and Powell [20] study short-term trading for a battery storage on real-time markets as a Markov Decision Process. Braeuer et al. [4], and more recently, Seifert et al. [34], present a European multi-market trading strategy for battery systems across three markets, trading at discrete time intervals and connecting the sequential reserve, day-ahead and intraday market decision stages. With a 4-hour time resolution, their ability to react to short-term changes in prices is very limited. Furthermore, the resulting smoothing of price spikes significantly reduces the profits a battery could extract from intraday markets. While co-optimizing storage assets across multiple markets (e.g., day-ahead, reserve, and intraday) is highly relevant in practice, integrating the continuous intraday market into multi-market frameworks currently requires heavy simplifications, such as artificially discretizing time into fixed decision stages. These simplifications severely distort the true profit potential of high-frequency trading. We focus exclusively on capturing the full complexity of the continuous intraday market because its rapidly growing volumes and high-frequency dynamics demand specialized, highly reactive algorithms. By isolating this market, we provide a highly accurate, fast-executing benchmark standalone model. We view understanding and accurately modeling this specific market as a critical foundational step before such algorithms can be expanded into broader multi-market environments without sacrificing intraday realism. Furthermore, sophisticated operation within the continuous intraday market is becoming increasingly critical for battery storage systems, as Europe's ancillary markets face saturation from the surging volume of installed battery capacity.

The myopic rolling intrinsic (RI) trading strategy is nearly optimal for gas storage operation [26]. The strategy naturally translates to other storage types, like batteries. Consequently, the RI is frequently used as a benchmark for more sophisticated strategies [e.g., 2,3].

Recent work by Semmelmann et al. [35] is most closely related to our work presenting an evaluation of RI trading for a battery on the continuous intraday market. However, their study relies on aggregated transaction data rather than the full limit order book, missing critical aspects of the market's microstructure, and does not account for battery degradation costs. While we adopt a simplified linear degradation model in this paper to maintain computational tractability at high execution frequencies, explicitly incorporating these costs into the optimization maps trading decisions more realistically and prevents the over-utilization of the asset.

Generally speaking, there are only a handful of papers that model the intraday market in its full complexity and implement trading strategies that take into account the complete information in the order book and, at the same time, do not artificially discretize time. The few examples of such papers known to the authors include Bertrand and Papavasiliou [2], Boukas et al. [3] and Kuppelwieser and Wozabal [23]. However, all of these papers are based on trading rules that are simple to execute. To the best of our knowledge, there is no single study of the rolling intrinsic policy or any other policy requiring a complex optimization for every single change in the limit order book while modeling detailed market behavior and rules.

This paper contributes to closing this gap by making the following main contributions:

1. We adapt the rolling intrinsic strategy to account for continuous trading on intraday markets for electricity, explicitly modeling every single order in the limit order book. The optimization is formulated as a mixed integer linear programming problem, taking into account all relevant aspects of the problem, including order placement costs, a simplified linear battery degradation model, and detailed market rules.
2. We approximate the exact MILP formulation using a dynamic programming approach, effectively trading a marginal loss in theoretical precision for a computational speedup of two to three orders of magnitude. As a result, our work enables the first high-frequency algorithmic trading strategy for storage assets, enabling realistic trading at every relevant update of the order book. Furthermore, we provide a detailed comparison, in terms of speed and profit, of our method against the MILP solution.
3. In a numerical study, we conduct a detailed backtest of our strategy in the German intraday market over a full year using an order-by-order traversal of historical order book data. The results show the revenue potential of a BESS operating in the continuous intraday market.

   The computational speed of our proposed algorithm is essential for such a detailed analysis, and our findings demonstrate that the speed of trading is critical to maximizing profits. Slower strategies, or generally less frequent trading, generate significantly lower profits by missing numerous trading opportunities due to their inability to make decisions at every relevant point in time. Our high-frequency strategy traverses the full year 2021 in around 86min, solving the intrinsic optimization approximately 24 million times (4.6 solves per millisecond), submitting around 30k orders at the exchange. Furthermore, we evaluate the robustness of our strategy under optimization and battery parameter settings, providing a deeper understanding of the method's performance and generalizability.
4. We parametrize the standard formulation of the RI and show that our simulation speed can be used to easily find optimal parameters, generating additional profits of 8.5 %, by slightly nudging the RI's trading behavior at no additional risk.
5. Lastly, we publish an easy-to-use Python package to run RI simulations over extended periods of time, given a set of battery and dynamic programming parameters. This will allow other researchers or industry professionals to realistically simulate high-frequency rolling intrinsic trading scenarios for a battery on the intraday market, thereby creating a strong benchmark that can be evaluated with minimal effort.

The remainder of this paper is organized as follows: Section 2 offers an overview of the European Power Market, setting the stage for our proposed trading strategy, detailed in Section 3. We present our results in Section 4, and discuss the details and implications of our findings. Finally, Section 5 concludes the paper and outlines directions for future research.

2. The European power market

The European short-term electricity markets are served by two primary competing power exchanges, Nord Pool and EPEX SPOT. While Nord Pool originated in the Nordic region, both exchanges now operate across wide areas of Europe, including the Nordic and Central Western European (CWE) regions.

The spot markets, which include day-ahead and intraday markets, are key components of the European Single Day-Ahead Coupling (SDAC) and Single Intraday Coupling (SIDC) initiatives. These initiatives have unified market operations across Europe, enabling bids submitted in one bidding zone to be seamlessly integrated into a shared order book and traded across other zones, provided there is sufficient cross-border transmission capacity. Price differences between bidding zones are thus primarily the result of constraints in interconnectors.

In this paper, we mainly focus on the German market, which is the largest electricity market in Europe and where trading is conducted mainly through the EPEX exchange. However, due to the ongoing convergence of the European electricity market designs, the proposed methods directly carry over to most other European markets.

The cascading range of future markets in Germany spans from long-term markets to day-ahead and intraday markets trading hourly and subhourly contracts. In continuous intraday trading, 60, 30 and 15 minute products are traded up to 5 minutes before delivery.

2.1. The intraday market for power

The primary objective of short-term power market designs is to minimize system imbalances. The intraday (ID) market plays a crucial role as the final opportunity for market participants to adjust their positions in response to unforeseen changes in production and demand, often driven by updated weather forecasts. Participants aim to minimize contractual deviations and reduce potential balancing costs charged by the transmission system operator (TSO). Significant price spreads between the day-ahead and intraday markets create strategic opportunities. Participants can withhold volumes from the day-ahead market to trade them intraday for additional profit [28,34].

Intraday market designs in Europe vary by country, although there has been significant progress toward a more unified ID market structure through the Single Intraday Coupling (SIDC) initiative, which includes both continuous and recently also auction-based trading products [10]. In this paper we focus on the former. While intraday auctions clear accumulated orders at specific, discrete times at a uniform price, the continuous market matches incoming orders instantaneously against the limit order book, as discussed in detail by Neuhoff et al. [30].

Unlike the day-ahead market, which clears via a single daily auction and requires only one discrete trading decision per day, the continuous intraday market operates through continuous matching and requires frequent, high-speed rebalancing. While the continuous ID market typically exhibits lower liquidity at any given point in time compared to the DA auction, it features significantly higher price volatility. This combination of unpredictability and continuous trading makes the ID market structurally more challenging to navigate, but it also creates highly rewarding arbitrage opportunities for fast-responding flexibility providers like battery storage systems.

In Germany, the intraday continuous market opens in the afternoon following the clearing of the day-ahead market. Trading begins at 3 p.m. on the day before delivery, allowing a maximum of 32 hourly or 128 quarter-hourly products to be traded at any given time. Trading for products in the shared SIDC order book closes 60 minutes before delivery, while products listed solely in the German order book trade until 30 minutes before delivery. At that point, the market splits into the four German TSO regions, in each of which trading continues until 5 minutes before delivery.

Limit orders submitted to the exchange are either immediately matched with corresponding buy or sell orders or, if unmatched, are stored in the limit order book [see 14]. A limit order is defined by its

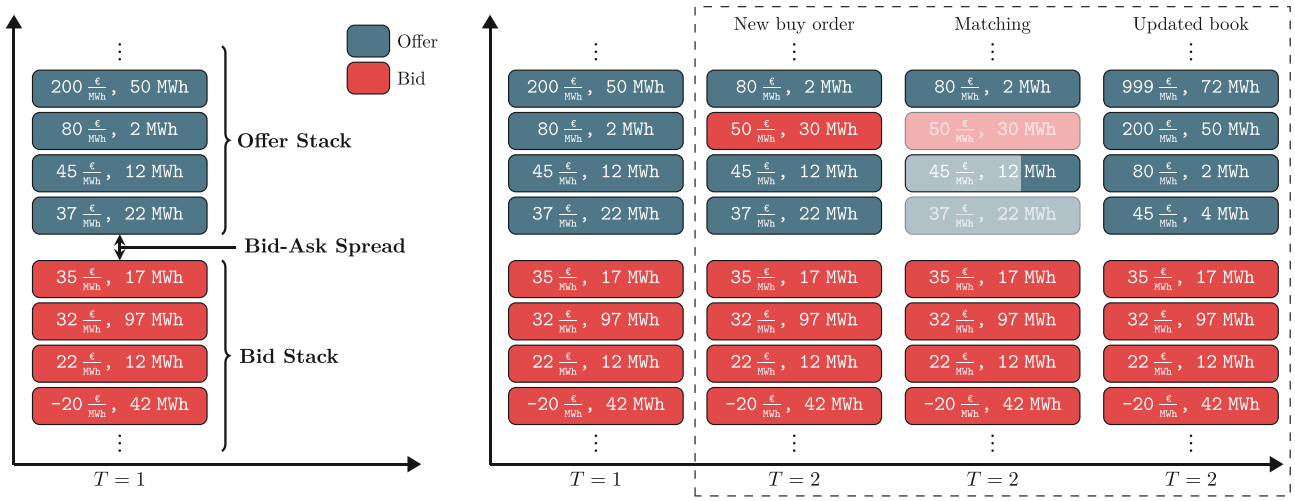


Fig. 1. An exemplary state of the LOB at $T = 1$ is presented in the left panel. The right panel depicts the clearing of a new hypothetical buy order: A new buy order with a price of 50 €/MWh, which is higher than the lowest ask price, is added to the book at $T = 2$. The quantity of this buy order (30 MWh) is then cleared against the cheapest possible offers until either the whole order is fulfilled (as is the case in the figure) or there are no offers with lower prices left. In this example, 22 MWh out of 30 MWh are cleared against the sell order with price of 37 €/MWh and the remaining 30 MWh − 22 MWh = 8 MWh are cleared against the sell order with price of 45 €/MWh. The remaining quantity of 12 MWh − 8 MWh = 4 MWh of the latter order stays in the order book. Note that the clearing is instantaneous, i.e., columns 2–4 in the right panel are purely illustrative and do not correspond to market states that can be observed by traders. Figure and caption adapted from Graf et al. [14].

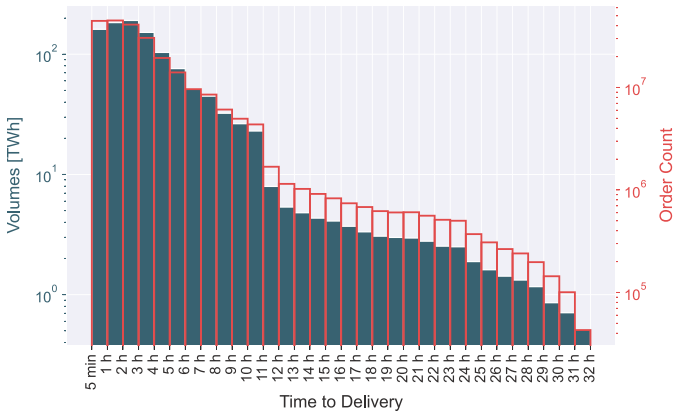


Fig. 2. The binned logarithmic distribution of time intervals between order submission and physical delivery for the German continuous intraday market in 2021. The teal bars illustrate the volume traded, while the red bars show the same data, counting the number of orders submitted. Both distributions exhibit an approximately exponential increase as the time to delivery shortens. Notably, there is a distinct gap between the 10-hour and 11-hour bins, where a significant increase in trading activity is observed.

price, volume, validity, and order type (buy/sell). Fig. 1 visualizes the process of a typical submission of a new limit order to the limit order book (LOB) of a product, i.e., delivery hour.

Continuous markets enable participants to continuously update and optimize their orders, allowing for real-time adjustments to evolving market conditions and reactions to forecast-updates. This flexibility makes them a crucial element in automated short-term trading strategies, where a rapid response to price changes is essential. In addition, continuous markets offer traders the opportunity to exploit short-lived market inefficiencies, enhancing their ability to capitalize on arbitrage opportunities throughout the trading day.

A key challenge of trading in the continuous ID market is the lack of liquidity several hours before delivery. This results in wide bid-ask spreads and high price volatility. An example of this liquidity disparity is shown in Fig. 2, which reveals that while a few hours before delivery typically thousands of orders are posted and cleared for a given product, this number drops to several hundred for products whose delivery is 10 hours or longer away and to a handful of orders at the start of trading.

As explained in more detail in [23], low liquidity far from delivery occurs because market participants lack precise information regarding their expected production and demand. As the delivery time approaches, weather forecasts and demand predictions become significantly more accurate. This increased certainty allows participants to take more precise balancing actions, which consequently drives up trading activity and liquidity.

Low liquidity creates difficulties for automated trading strategies that aim to capitalize on price differences throughout the day. For instance, evening products traded during the morning hours often fail to reflect prevailing market prices accurately, as they are largely influenced by market makers rather than real supply and demand dynamics. This issue of liquidity is the main disadvantage of the continuous market, compared to its auction alternatives. It reduces the efficiency of price discovery and also increases traders' risk, as they may experience suboptimal pricing and reduced profitability when locking in positions early, as discussed more in-depth in e.g., [14].

Finally, Table 1 provides a concise summary of key statistics for the Central-Western European (CWE) Spot Market over recent years. The continuous market stands out as the dominant intraday trading platform. Specifically, the German continuous intraday market in 2021, which serves as the data source for this study, recorded 240.4 million order submissions (including order changes). This translates to an average of 7.6 orders placed per second. The sheer volume and frequency of transactions underscore the necessity of accurately modeling and responding to high-frequency market dynamics, to realistically simulate any trading strategy on this market. This calls for extremely fast and reactive trading. In the next section, we will propose an algorithm that can easily keep pace with the market's speed.

3. An efficient rolling intrinsic policy for continuous markets

In this section, we describe the intrinsic problem for continuous markets, discuss the rolling intrinsic policy as an improvement of the

Table 1
Combined EPEX Spot trading volumes in TWh for the CWE region (AT, BE, DE/LU, FR, NL) taken from the EPEX Spot annual reports [9]. The percentage-change to the respective last year is given in brackets. Clearly, intraday trading is gaining relevance compared to DA auctions, while the ID auction market is slowly picking up in activity.

	2023	2022	2021	2020	2019	2018	2017
Day-Ahead	(+16.9%) 419.3	(−8.6%) 358.8	(−4.5%) 392.7	(−8.4%) 411.1	(+3.1%) 448.8	(+5.3%) 435.2	413.2
ID-Auction	(+9.0%) 9.41	(−1.9%) 8.63	(+0.0%) 8.8	(+20.5%) 8.8	(+7.4%) 7.3	(+30.8%) 6.8	5.2
ID-Continuous	(+29.9%) 119.8	(+8.3%) 92.2	(+10.1%) 85.1	(+23.5%) 77.3	(+11.4%) 62.6	(+11.3%) 56.2	50.5

intrinsic strategy, and present an efficient implementation as a sequence of dynamic programming problems.

3.1. The intrinsic problem

The intrinsic strategy is a simple policy to maximize the rewards of a storage system, which is widely used in the literature on gas storage and in the industry to optimize storage assets [e.g., 15,16,24–26]. In the extant literature, the intrinsic problem is set in a perfectly liquid futures market, trading $T \in \mathbb{N}$ futures of a commodity that can be stored and which trade for prices P_t. To avoid overcomplicating notation, we assume here that futures contracts are for delivery in consecutive time periods $[0,1), [1,2), \dots, [T-1, T)$ and that there are no overlapping contracts. To fix ideas, one can imagine a battery storage system trading futures for delivery in the 24 hours of the next day.

The storage operator decides on traded quantities q_t (MWh) of these contracts such that volumes bought are transferred into the storage, while sold volumes are physically fulfilled from the storage. In particular, in the simplest case, the intrinsic optimization boils down to the following linear problem of optimizing the storage level $(s_t)_{t=1}^T$ (MWh) and the future positions $(f_t)_{t=1}^T$ (MWh) in the following way:

$$\begin{aligned} \max_{q_t} \quad & \sum_{t=1}^{T} -P_t q_t \\ \text{s.t.} \quad & f_t = f_t^0 + q_t, && \forall t = 1, \dots, T \\ & f_t \in [\underline{f}, \bar{f}], && \forall t = 1, \dots, T \\ & s_t = s_{t-1} + f_t, && \forall t = 1, \dots, T \\ & s_t \in [0, \bar{s}], && \forall t = 1, \dots, T, \end{aligned} \tag{1}$$

where s_0 and $(f_t^0)_{t=1}^T$ are the initial storage level and the initial future positions, respectively, which are data for the problem. In the above problem, positive q_t represents buying decisions while negative ones represent selling decisions. Furthermore, $\bar{s}$, $\underline{f}$, and $\bar{f}$ are physical bounds for the maximum storage level (MWh) as well as for withdrawal and injection (MWh). The respective constraints ensure that the position is *physical* in the sense that traded quantities can be accommodated in the storage without violating the limits on injection and withdrawal.

The above linear problem is the simplest form of the intrinsic, which assumes a perfectly efficient storage, liquid markets, and no cost for operating the storage. Furthermore, the problem is myopic in the sense that the decision maker uses all the flexibility of the storage instantly without anticipating future changes in prices, which would make it more advantageous to wait before trading immediately. The myopic nature of the intrinsic policy makes it suboptimal and motivates its name originating from the *intrinsic value* of a financial option contract, i.e., the value an option has when it is immediately exercised, not taking into account the value of waiting (the extrinsic value). Problem (1) is the version of the intrinsic problem prevailing in the literature.

Next we will extend (1) to a policy for a BESS trading on continuous intraday power markets, explicitly taking into account order book information, trading costs, and a simple linear model for efficiency and degradation losses of batteries. In order to do so, we reuse Eq. (1)'s notation and introduce a set of futures contracts $\mathcal{T}_{t^*}$ at time t^*, for delivery of electricity in non-overlapping time periods in the future. The time when the earliest contract goes into delivery is denoted by t_0. Note that the formulation could easily be extended to futures contracts that overlap, such as quarter-hourly and hourly contracts. For simplicity, in the following we only focus on the intrinsic at a singular time t^*, denoting $\mathcal{T}_{t^*}$ simply as $\mathcal{T}$.

For each contract $t \in \mathcal{T}$, there is an order book $\mathcal{O}_t = \mathcal{O}_t^+ \cup \mathcal{O}_t^-$ with $\mathcal{O}_t^+$ containing all the currently active asks and $\mathcal{O}_t^-$ containing all the currently active bids for contract t. For each of these orders $i \in \mathcal{O}_t$, we define a limit price P_i and a quantity $Q_i > 0$. Within this framework, the optimization evaluates the exact historical limit order book to determine which orders to match. By design, the objective function naturally prioritizes the most competitive prices, always taking the lowest available asks and highest available bids, since selecting a less favorable price for a given volume would yield a mathematically suboptimal result. Furthermore, we introduce efficiency factors η^+, $\eta^- \in (0, 1]$ for charging and discharging, respectively, as well as a volume-based trading cost ν_{trade} for matching orders on the market.

Batteries have a limited life. Roughly speaking, there are two effects causing battery degradation: calendar age, i.e., the loss of capacity with time even without active use of the battery, and cycling, i.e., the wear and tear caused by injection and withdrawal of energy. Since we are looking at short-term planning, we focus on the latter and model a symmetric linear cycling (i.e., degradation) cost ν_{deg} in €/MWh of injected/withdrawn energy. The intuition behind this choice is that cells will have to be replaced after a certain number of cycles, which induces a replacement cost. With this information, we calculate a proportional per cycle and ultimately a per MWh cost for using the battery. We note that actual degradation depends not only on the amount of cycled energy but also on the depth of discharge, temperature, and several other factors [see e.g., 17,38]. However, for the sake of simplicity, we focus on the amount of cycled energy, which is the most important factor. Combining the cost of trading and the cost of degradation, we define $\nu = \nu_{\text{trade}} + \nu_{\text{deg}}$ as the combined variable cost in units of €/MWh. By acting as a friction term in the objective function, this combined cost ensures that the battery only executes trades where the market price spread exceeds the cost of physical wear and tear. Consequently, the degradation model actively prevents marginal trades, which reduces the total executed trading volume and theoretical profits, but yields a realistic and economically viable dispatch schedule.

Until this point, all discussed extensions of (1) can still be modeled as a linear program (LP). However, there are two aspects of intraday trading that introduce non-linearities. Firstly, trade quantities on electricity exchanges are not continuous variables but multiples of the minimum trading unit u, which implies that they are of the form $q_i = k_i u$ with $k_i \in \mathbb{N}_0$. Secondly, efficiencies smaller than 1 coupled with negative prices make it profitable to spend energy by simultaneously charging and discharging, which is physically not possible for a battery. More specifically, for a single product t and orders $i \in \mathcal{O}_t^+$, $j \in \mathcal{O}_t^-$, buying a quantity q using order i and immediately re-selling the power $\eta^-(\eta^+ q)$ that remains after accounting for efficiency losses using order j yields a positive profit if

$$-qP_i + \eta^-(\eta^+ q)P_j > 0 \Leftrightarrow (\eta^-\eta^+)P_j > P_i. \tag{2}$$

Since $P_i > P_j$ by definition, this can only happen if both prices are negative. However, because it is physically not possible to charge and discharge the battery at the same time, such trades must be prevented.

Algorithm 1 The rolling intrinsic.

Input: Initial storage s_0, start-time t_{start} and stop-time of simulation t_{end}, limit order book information $\mathcal{T}_{t^*}$ $\forall t^* \in [t_{start}, t_{end}]$, solve-frequency of intrinsic, other parameters $(\bar{f}, \underline{f}, \nu, \dots)$
Output: Final battery schedule during simulation runtime, cumulated trade profits
for time instant $t^* \in [t_{start}, t_{end}]$ **do**
  Solve the intrinsic according to Eq. (3)
  Update all positions $f_t, \forall t \in \mathcal{T}_{t^*}$
  Log profits/losses from trading
end for

With these preparations, we can write the modified intrinsic as the following mixed integer linear problem (MILP):

$$\max_{f_t, s_t, \alpha_t, q_i, k_i} \quad \sum_{t\in\mathcal{T}} \left(\sum_{i\in\mathcal{O}_t^-} (P_i - \nu) q_i - \sum_{i\in\mathcal{O}_t^+} (P_i + \nu) q_i \right) \tag{3}$$

$$\begin{aligned}
\text{s.t.} \quad & 0 \le q_i \le Q_i, && \forall i \in \mathcal{O}_t, \forall t \in \mathcal{T} \\
& q_i = k_i u, && \forall i \in \mathcal{O}_t, \forall t \in \mathcal{T} \\
& f_t^+ = \sum_{i\in\mathcal{O}_t^+} q_i, && \forall i \in \mathcal{O}_t^+, \forall t \in \mathcal{T} \\
& f_t^- = \sum_{i\in\mathcal{O}_t^-} q_i, && \forall i \in \mathcal{O}_t^-, \forall t \in \mathcal{T} \\
& f_t = f_t^0 + f_t^+ - f_t^-, && \forall t \in \mathcal{T} \\
& f_t \in [\underline{f}, \bar{f}], && \forall t \in \mathcal{T} \\
& f_t = i_t - w_t, && \forall t \in \mathcal{T} \\
& i_t \in [0, \alpha_t \bar{f}], && \forall t \in \mathcal{T} \\
& w_t \in [0, -(1-\alpha_t)\underline{f}], && \forall t \in \mathcal{T} \\
& s_t = s_{t-1} + \eta^+ i_t - \frac{1}{\eta^-} w_t, && \forall t \in \mathcal{T} \\
& s_t \in [0, \bar{s}], && \forall t \in \mathcal{T} \\
& \alpha_t \in \{0,1\},\ k_i \in \mathbb{N},\ \nu \in \mathbb{R}_{\ge 0} && \forall i \in \mathcal{O}_t, \forall t \in \mathcal{T}
\end{aligned}$$

where the α_t are binary variables which are equal to 1 in case the decision for contract t is to buy and 0 if the decision is to sell, and i_t (f_t^+) and w_t (f_t^-) are the combined buy and sell positions per product including (excluding) previous positions.

Note that, due to the above discussion, strictly speaking, the binary variables α_t are only required for hours t where both $\mathcal{O}_t^+$ and $\mathcal{O}_t^-$ contain orders with negative prices. Since u is usually rather small and negative prices do not occur too often, typically linear relaxations of the above problems are good and yield near-optimal solutions in the root node of the branch-and-bound trees, making the problem a comparably easy MILP in most realistic cases.

3.2. The rolling intrinsic

The rolling intrinsic policy (RI), as outlined in Algorithm 1 was originally introduced by Gray and Khandelwal [16] and is a dynamic extension of the intrinsic value discussed in the previous section. Starting from an initial storage state s_0 and market positions $(f_t)_{t\in\mathcal{T}}$ acquired in previous periods, the RI repeatedly checks for opportunities for profitable rebalancing by re-running the intrinsic policy. Although the resulting decisions are still myopic, the RI clearly represents an improvement over the static intrinsic policy which does not adapt positions at all. Furthermore, the myopic nature of the decisions has the advantage that the RI does not speculate but only enters immediately profitable positions and therefore does not run the risk of accumulating losses. For these reasons and because of its conceptual simplicity and relatively low computational cost, the RI has gained widespread industry adoption.

While the policy remains strictly myopic, its computational simplicity is precisely what allows it to be evaluated at extremely high frequencies. In a continuous market, the ability to instantly capture fleeting, short-lived price spreads through rapid execution often yields an operationally stronger and more profitable strategy than theoretically more sophisticated but computationally slow and practically unimplementable models.

3.3. A dynamic programming formulation of the rolling intrinsic policy

In this section, we describe how to reformulate problem (3) into a dynamic programming formulation that can be solved by an efficient implementation of the dynamic programming algorithm. The aim of the reformulation is to make faster decisions to keep up with the pace of the intraday market, where typically many orders arrive per second, requiring constant rapid reevaluation of one's position.

The main idea of the reformulation is to assign each product $t \in \mathcal{T}$ a stage in which the decisions to accept orders $i \in \mathcal{O}_t$ are taken while the storage level is the state variable that connects the stages. Hence, in the language of dynamic programming, the decisions q_i, k_i, α_t, and f_t are the actions, and the resulting storage state s_t is the argument of the value function of the next stage. The solution of the intrinsic problem itself remains a single-step optimization, executed at a specific point in time. To solve this high-dimensional decision problem efficiently, we break up the logic of the underlying mixed-integer linear program into individual products. Mathematically, each delivery hour is mapped to a distinct "stage" of the DP. Note that this construction is artificial, as the DP stages do not represent different points in time when sequential decisions are made, but merely different parts of a single trading decision that is taken all at once.

We assume without loss of generality that $\mathcal{T}_t = \{1, \dots, T\}$ set $V_{T+1} \equiv 0$ and for $t = 1, \dots, T$ define the value function as

$$V_t(s_{t-1}) = \begin{cases}
\max\limits_{f_t, s_t, \alpha_t, q_i, k_i} & \sum\limits_{i\in\mathcal{O}_t^-} (P_i - \nu) q_i - \sum\limits_{i\in\mathcal{O}_t^+} (P_i + \nu) q_i + V_{t+1}(s_t) & \\
\text{s.t.} & 0 \le q_i \le Q_i, & \forall i \in \mathcal{O}_t \\
& q_i = k_i u, & \forall i \in \mathcal{O}_t \\
& f_t^+ = \sum_{i\in\mathcal{O}_t^+} q_i, & \forall i \in \mathcal{O}_t^+ \\
& f_t^- = \sum_{i\in\mathcal{O}_t^-} q_i, & \forall i \in \mathcal{O}_t^- \\
& f_t = f_t^0 + f_t^+ - f_t^- & \\
& f_t \in [\underline{f}, \bar{f}] & \\
& i_t \in [0, \alpha_t \bar{f}] & \\
& w_t \in [0, -(1-\alpha_t)\bar{f}] & \\
& s_t = s_{t-1} + \eta^+ i_t - \frac{1}{\eta^-} w_t & \\
& s_t \in [0, \bar{s}] & \\
& \alpha_t \in \{0,1\},\ k_t \in \mathbb{N}, \nu \in \mathbb{R}_{\ge 0}. &
\end{cases} \tag{4}$$

This formulation might not seem very natural, since in reality the decisions in all stages of the above problem happen simultaneously instead of sequentially as is usually the case in dynamic programs. Because of this, existing studies and industry practice almost exclusively rely on monolithic MILP or LP formulations, which offer the most intuitive and direct mathematical mapping of the battery's physical constraints. However, reformulating this simultaneous problem into a multi-stage dynamic program is the key to unlocking significantly faster computation times. As we will demonstrate in Section 4, this novel perspective allows us to solve the intrinsic optimization up to three orders of magnitude faster than a standard MILP, executing several solves per millisecond. This speedup is not merely a convenience; it is a strict prerequisite for realistically tracking and reacting to the millisecond-pace of the continuous limit order book.

We propose solving the problem by applying the dynamic programming algorithm, discretizing the actions f_t into steps $\kappa \cdot u$ ($\kappa \in \mathbb{N}$) apart, where u is the minimal tradable quantity. Furthermore, we introduce a function π_t that encodes the current order book information and returns the cost or revenue of buying or selling on the intraday market, taking into account trading costs. Additionally, we define a state transition

function

$$S(s_{t-1}, f_t) = \begin{cases} s_{t-1} + \eta^+ f_t, & \text{if } f_t > 0 \\ s_{t-1} + \frac{f_t}{\eta^-}, & \text{otherwise.} \end{cases} \tag{5}$$

With these preparations, we can now rewrite problem (4) as

$$V_t(s_{t-1}) = \begin{cases} \max\limits_{k_t} & \pi_t(f_t^0 - f_t) + V_{t+1}(S(s_{t-1}, f_t)) \\ \text{s.t.} & f_t = f_t^0 + k_t u \\ & k_t \in \left[-\frac{\eta^- s_{t-1} + f_t^0}{u}, \frac{\bar{s} - s_{t-1}}{\eta^+ u} - \frac{f_t^0}{u} \right] \cap \left[\frac{\underline{f} - f_t^0}{u}, \frac{\bar{f} - f_t^0}{u} \right] \cap \kappa\mathbb{Z}, \end{cases} \tag{6}$$

where the bounds in the first constraint on k_t enforce the storage energy limits, while the second interval enforces the power limits. Given V_{t+1}, solving (6) boils down to evaluating the objective for all feasible k_t, i.e., to a small number of arithmetic operations.

Note that the above reformulation of (4) to (6) automatically takes care of the issue of simultaneously accepting buy and sell orders in the case of negative prices, and, by fixing the trading step-size u to a multiple of the minimum bid-size, naturally yields feasible position sizes.

To solve (6), we need to know the value functions V_{t+1}. In line with the usual dynamic programming algorithm, we calculate V_t going backward in time: V_{T+1} is the known boundary condition, which makes it possible to calculate V_T in any given state. However, due to inefficiencies $\eta^\pm < 1$, the trades that live on a finite grid translate to storage levels that do not sit on a regular grid, especially after several rounds of injections and withdrawals.

Therefore, unlike classic applications of the dynamic programming algorithm, we do not define V_T to have a finite domain, but instead define it on the whole interval of possible storage states $[0, \bar{s}]$. In order to obtain the value function, we discretize the possible storage states $[0, \bar{s}]$ to a finite grid $G \subset [0, \bar{s}]$ with $|G| = m$. We then approximate the value function V_T by evaluating it on the grid and linearly interpolating for points $s \notin G$. Thus, the parameter m represents the number of discrete grid points we use to model the possible state of charge of the battery within our dynamic programming approach. A smaller storage discretization, such as $m = 11$, means we evaluate the value function at fewer points, resulting in faster computation times but a slightly less accurate approximation of the optimal strategy.

In particular, we denote the resulting approximations by $\tilde{V}_{t+1}$ and describe our linear approximation between the neighboring discretized grid points $s_i, s_{i+1} \in G$ as follows:

$$\tilde{V}_{t+1}(s_t) = \frac{m(s_{i+1} - s_t)}{\bar{s}} V_{t+1}(s_i) + \frac{m(s_t - s_i)}{\bar{s}} V_{t+1}(s_{i+1}), \tag{7}$$

for $s_i \le s_t \le s_{i+1}$. More sophisticated interpolations, taking into account, for example, the curvature of the value function (given by the buy and sell price stacks), could potentially enhance the precision of our optimization. However, since such approaches would be computationally more costly and the obvious choices, such as quadratic interpolations or cubic splines, did not show any improvements in our numerical experiments, we remain with the linear approximation above.

Once we have obtained $\tilde{V}_T$, we can solve the problem that defines V_{T-1} with V_T replaced by $\tilde{V}_T$. We then repeat the process until we have obtained the approximations $\tilde{V}_t$ of all the value functions. Note that, in principle, G can be chosen independently for every t and can also change between various solutions of the intrinsic problem. For the sake of simplicity, we keep the grid constant, i.e., $G_{t^*,t} \equiv G$. Thus our DP's computational complexity grows linearly as $\mathcal{O}(\frac{|T| \cdot |G|}{\kappa})$.

Algorithm 2 displays an overview of one such dynamic programming solution run of the intrinsic problem at time t^*, where $\mathcal{X}_t$ is the set of all possible actions k_t as defined in Eq. (6).

The approximation of the value functions $\tilde{V}_t$ introduces an error relative to the exact MILP formulation of the problem. In the following proposition, we show that the domain of the true functions V_t is not actually continuous and that consequently there is a grid G, which makes the approximation and thereby the DP formulation of the problem exact.

Algorithm 2 DP at runtime t^*.

Input: f_t^0, $\mathcal{O}_t$ for all tradable products $t \in \mathcal{T}$, other parameters $(\kappa, s_0, \nu, G, \dots)$
Output: Updated market positions f_t
Set $V_{T+1} \equiv 0$.

Backwards Pass
for $t = T, \dots, t_0$ **do**
  for each state $s \in G$ **do**
    $V_t(s) \leftarrow \max_{k_t \in \mathcal{X}_t(s, f_t^0)} \{\pi_t(k_t u) + \tilde{V}_{t+1}(S(s, k_t u))\}$
  end for
end for

Forward Pass
for $t = t_0, \dots, T$ **do**
  $\hat{f}_t \leftarrow arg\,max_{k_t \in \mathcal{X}_t(s, f_t^0)} \{\pi_t(k_t u) + \tilde{V}_{t+1}(S(s_{t-1}, k_t u))\}$
  $f_t^0 \leftarrow \hat{f}_t$
  $s_t \leftarrow S(s_{t-1}, \hat{f}_t)$
end for

Proposition 1. *For all $t \in \{1, \dots, T\}$, there is a finite grid G_t that contains all possible storage states at the beginning of period t. Choosing G to approximate the value functions alleviates the need for interpolation and makes the approximation exact.*

Proof. Suppose that the initial state of charge (SoC) is s_0, we set $\kappa = 1$ and look at the simplest case of $f_t^0 = 0$ for all $t \in \mathcal{T}$. The possible storage states at the beginning of stage $t_0 + 1$ are

$$G_{t_0+1} = \left\{ s_0 + u k_1^+ \eta^+ - \frac{u k_1^-}{\eta^-} : k_1^\pm \in \mathcal{K}_1^\pm \right\}, \tag{8}$$

and $\mathcal{K}_1^+ = \left\{ k \in \mathbb{N}_0 : k \le u^{-1} \min\left(\frac{\bar{s} - s_0}{\eta^+}, \bar{f} \right) \right\}$ and $\mathcal{K}_1^- = \left\{ k \in \mathbb{N}_0 : k \le u^{-1} \min\left(s_0 \eta^-, -\underline{f} \right) \right\}$.

Defining the sets G_t and $\mathcal{K}_t^\pm$ going forward in time, for period $t > t_0$ we thus have, given that $\mathcal{K}_{t-1}^\pm$ is already known,

$$G_t = \left\{ s_t = s_0 + u k_t^+ \eta^+ - \frac{u k_t^-}{\eta^-} : k_t^\pm \in \mathcal{K}_t^\pm, 0 \le s_t \le \bar{s} \right\} \tag{9}$$

with

$$\mathcal{K}_t^+ = \left\{ k \in \mathbb{N}_0 : k \le \frac{(\bar{s} - s_0) + u \max(\mathcal{K}_{t-1}^-)(\eta^-)^{-1}}{u \eta^+} \right\} \tag{10}$$

$$\mathcal{K}_t^- = \left\{ k \in \mathbb{N}_0 : k \le u^{-1} \eta^- \left(s_0 + \eta^+ u \max(\mathcal{K}_{t-1}^+) \right) \right\}. \tag{11}$$

Note that the way G_t and $\mathcal{K}_t^\pm$ are defined takes care of repeated charging and discharging in periods $t_0, \dots, t$ and therefore covers all possible storage states. □

Choosing the finite grids described above ensures that the DP solution is exact and equivalent to the MILP formulation, as the value function only needs to be evaluated at points in the grid and there is no need for approximation. This comes at the cost of computational efficiency, as choosing G to be exact, would yield a very fine grid for cases where there are many tradable products $\mathcal{T}$ and $\eta^\pm < 1$. The formulation presented above adds to the extant literature in two important ways:

1. **Full LOB Granularity:** Unlike standard rolling intrinsic implementations that rely on aggregated price curves, transaction data,

or simplified mid-market prices, our approach explicitly models the *full limit order book dynamics*. This allows the policy to capture liquidity constraints and depth-dependent pricing exact to the historic microstructure level.
2. **High-Frequency Feasibility via DP:** The reformulation from a generic MILP to a specialized dynamic programming scheme reduces the computational burden by orders of magnitude. This speedup is not merely an efficiency gain, it is a structural prerequisite that enables the backtesting of a strategy that re-optimizes on a millisecond timescale over a full year of market tick-data, a task that would be computationally intractable with standard solvers.

3.4. From optimization to execution

The optimal position changes $\hat{f}_t$ derived from the forward pass must be translated into actionable limit orders. To mitigate execution risk in a fast-moving market, we submit these quantities as fill-or-kill orders. This ensures that the strategy only trades a position if the full desired volume is available at the expected price, preventing partial fills that could lead to suboptimal inventory states. However, due to the latency Δt between the optimization start and order arrival at the exchange, market conditions may shift, potentially leaving orders unmatched or resulting in positions that, while optimal at t^*, are no longer physically feasible given the market's state at $t^* + \Delta t$. In such cases, the rolling nature of the policy corrects these deviations in the subsequent optimization step, naturally steering the system back toward feasible, profitable operation.

Furthermore, it is important to emphasize that while our algorithmic trading strategy operates at a high financial frequency, this does not strictly equate to high-frequency physical cycling of the battery. Frequent re-evaluations of the market state primarily allow the algorithm to adjust the timing of financial positions as the order book evolves. However, re-optimizing a myopic schedule in a volatile continuous intraday (CID) market can influence the physical dispatch, and thus the wear-and-tear of the battery, in three primary ways:

1. **Shifting volume:** The algorithm might shift the traded electricity between hours (e.g., moving delivery to a neighboring hour) to capture a better price. This keeps the total physical volume constant, though it requires the original positions to be unwound.
2. **Increasing volume:** The optimization may encounter new, independent price spreads, not affecting previous positions. Capitalizing on these new spreads layers additional charge and discharge actions onto the existing schedule, strictly increasing the battery's physical cycles.
3. **Reducing volume:** If market prices move such that an originally profitable spread collapses or inverts, the optimal action may be to only unwind the existing positions, thereby reducing the scheduled physical cycles.

In practice, any combination of these effects can occur for a given optimization instance. Depending on market conditions, total physical cycles might therefore increase or decrease as a result of high-frequency re-optimization. Consequently, our combined linear degradation and trading costs naturally act as necessary friction terms against weak market signals. This prevents the algorithm from constantly chasing marginal price movements, thereby regulating the schedule revisions and limiting the battery's actual physical wear and tear.

By contrast, alternative reduced-order market formulations, such as volume-weighted average prices, are computationally lighter and allow for richer battery models, as they frequently overestimate revenues by ignoring strict limit order book liquidity constraints. By operating directly on the limit order book, our framework prevents the policy from attempting to clear energy volumes that do not physically exist at the expected price. Furthermore, our strategy provides a clear, directly implementable policy on how and when the battery's orders should be submitted to the market and how one must dynamically correct missed positions to avoid physical imbalances.

4. Results

In this section, we present the results of a numerical case study. In Section 4.1, we discuss the setting and implementation of the case study. In Section 4.2, we measure the precision of our DP solution and compare the runtimes with the exact MILP formulation, arguing that the DP leads to runtimes that are orders of magnitude faster. In Section 4.3, we present the results of our policy for a whole year of trading and measure the impact of trading speed on performance, while in Section 4.4, we examine how changes in storage and optimization parameters affect the results.

4.1. Setting and mode of comparison

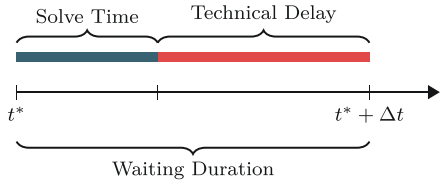
Unless otherwise stated, for our case study, we use a BESS with an energy-to-power ratio (duration) of 1h, set $\bar{f} = -\underline{f} = 10$ MWh (i.e., a power of 10 MW) and $\bar{s} = 10$ MWh, and assume efficiency losses of $\eta^+ = \eta^- = 0.95$, i.e., a round trip efficiency of $\eta^+\eta^- \approx 0.9$ similar to Cole and Karmakar [7] and Regelleistung Online [31].

We use linear degradation costs of $\nu_{\text{deg}} = 4$ €/MWh, i.e., a round-trip degradation cost of 8 €/MWh for every MWh of electricity sold by the battery. We motivate this value by Goldman Sachs [13] who project the undiscounted cost of battery replacement in 2030 to be around 60k €/ MWh, a 10-year battery warranty, and approximately 2 charging cycles per day. Dividing the replacement costs by cycles gives the cost of one full cycle, which allows us to infer the induced €/MWh linear degradation cost: $8 \approx 60k/(2 \cdot 365 \cdot 10)$.

Furthermore, we choose $\nu_{\text{trade}} = 0.09$ €/MWh, reflecting the current EPEX volume-based trading fees for each matched order. In total, the variable cost $\nu = \nu_{\text{trade}} + \nu_{\text{deg}}$ thus takes the value of $\nu = 4.09$ €/MWh. This variable, symmetrically applied, cost ν is always present in the optimization, while ν_{trade} is deducted from profits after each transaction. The degradation cost ν_{deg} is subtracted only once, after each product's final schedule is set at gate closure.

All results are generated using historical 2021 order book data from the EPEX Spot continuous intraday market for Germany, focusing on hourly products. It is not possible ex post to simulate iceberg orders correctly given the EPEX dataset, so we omit them in our simulations. The exchange allows for price increments of 0.01 € and volume increments at 0.1 MWh, defining the minimum action size on the market as $u = 0.1$ MWh. The discretization of the dynamic programming action is defined at the most accurate minimum value of $\kappa = 1$ throughout.

Clearly, in real-life trading, there will always be some delay between two consecutive optimizations. We explicitly model this delay Δt and conceptually split it into two components: the *solve-time* and the *technical delay*:



The solve time represents the time required to find an optimal solution for a given intrinsic problem. The technical delay accounts for various operational latencies, such as communication delays with the exchange, delays caused by internal databases and trading systems, and other post-solve delays before the submitted orders are matched on the exchange. If the combined delay exceeds 1 ms (the tick duration of the EPEX Spot market), this has two effects on trading: Firstly, there is the risk that an order placed by our policy may not be matched since its intended counterpart either expired or was cleared against other market participants' orders in the meantime. We deal with this by submitting all-or-none orders that expire if they are not immediately cleared, which

Table 2
Baseline simulation parameters for the BESS, the intraday market environment and the DP.

Parameter	Symbol	Value	Unit
Battery Storage Specs			
Energy Capacity	$\bar{s}$	10	MWh
Battery Duration	–	1	h
Efficiency (Charge/Discharge)	η^+, η^-	0.95	–
Costs & CID Market			
Trading Fee	v_{trade}	0.09	€/MWh
Degradation Cost (Linear)	v_{deg}	4.00	€/MWh
Min. Tradable Unit	u	0.1	MWh
Simulated Year	–	2021	–
Optimization Settings			
Action Discretization	κ	1	–
Storage Discretization Size	$\lvert G \rvert$	11	–

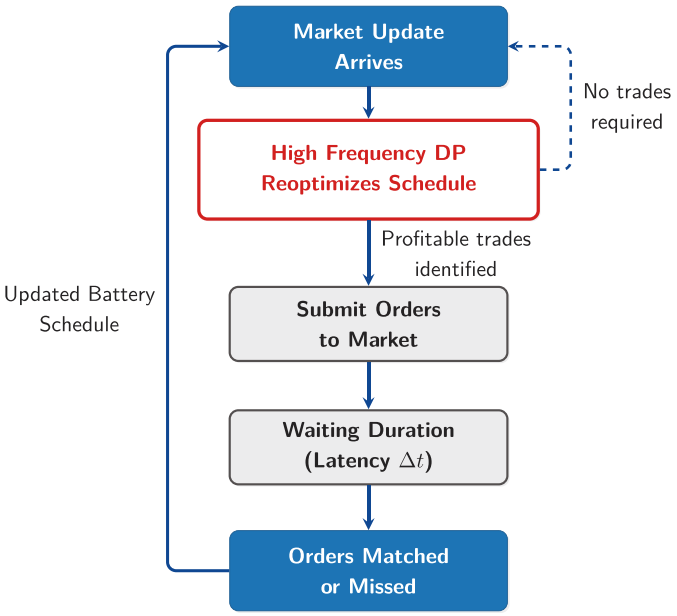


Fig. 3. Flowchart of our method's execution logic.

might lead to non-physical positions, i.e., market positions that cannot be matched by a feasible schedule of injections and withdrawals. In such a situation, further trades are required to correct the position so that the final battery schedule becomes physically feasible. However, in the way the intrinsic problem is set up, such positions would be corrected in the next call of the intrinsic problem, forcing the RI to return to a physical position as quickly as possible.

The second issue arising from delays is that the policy might not be able to react to every single order book update. In particular, we assume that the next intrinsic is solved only at $t^* + \Delta t$ observing the updated order book state at this time. Hence, if there are updates to the order book in the time interval $(t^*, t^* + \Delta t)$, the policy will not be able to react immediately to it. We give a short overview of our chosen standard parameters in Table 2, and provide a graphic overview of the simplified execution logic in Fig. 3.

Finally, we remark that when backtesting, we correctly account for changes in the order book caused by our actions. This in particular means that EPEX orders cleared against the orders of our policy are available for a shorter duration than in the actual history, and consequently that subsequent clearing of orders by other market participants changes. This introduces a strong path-dependency of varying trading policies.

Although we take great care to model the market correctly, a limitation of the backtesting experiments is that, by the very nature of our analysis and the available data, we cannot take into account the effect that the orders placed by the strategy would have had on the behavior of other market participants. However, this drawback is inherent in the idea of backtesting and cannot be easily corrected.

All simulations were implemented in C++ and run on a single AMD EPYC 9654 CPU, with the MILP solved using Gurobi v11.0.3 for C++ [18]. Due to the simple nature of each single intrinsic optimization, solving the MILP on multiple threads does not lead to faster simulation times. Additionally, to provide a fair comparison, we set our MILP solver's time-limit to 2 s (MIP gap 10%) and warm-start each optimization with the null-action, i.e., the solution of not changing the current battery schedule. We publish our code in the Python package *BitePy*, which allows users to easily run our DP simulations over a Python interface. See Appendix A for details.

4.2. Comparison between MILP and DP rolling intrinsic

To evaluate the effectiveness of solving the intrinsic rolling problem using the DP approach instead of the MILP approach, we start by comparing the two methods. The MILP consistently provides the exact solution to the optimization problem (3), while the DP solution generally slightly diverges. Note that, due to the RI's myopic nature and the path dependence of differing solutions, following the exact solution does not necessarily guarantee larger profits. For this reason, although inexact in solving single intrinsic problems, the DP-based rolling intrinsic strategy can be more profitable than the MILP formulation, as evidenced by the February and April results in Table 3.

We choose the most clinical setting for our comparison, by artificially fixing the solve time and technical delay to 0 ms for both approaches. This results in a policy that ignores technical delays and solves the intrinsic optimization for each relevant order book update in such a way that the resulting all-or-none limit orders are matched instantaneously with the exchange after each solve.

Table 3 reports this comparison for the first week of February, April, July, and October 2021. All three DP methods use a constant equidistant grid G, with storage discretizations $m = 101$, $m = 51$ and $m = 11$. In summary, we can say that even without delays, the DP achieves comparable rewards to MILP and significantly outperforms it in terms of runtime by a factor of 100–1000 for the coarser storage discretization $m = 11$. While the DP simulation times remain constant, MILP simulation times vary, even though the number of intrinsic solves is almost equal, as the complexity of the MILP problem is highly influenced by current market conditions. Interestingly, a finer storage discretization does not provide a clear advantage over coarser discretizations in this experiment, but shows an increase in profits for our yearly simulation results, as evidenced by the following sections.

Bootstrapping the test results reveals no significant difference ($p \approx 0.66$) in the hourly (per-product) rewards generated by the MILP compared to our DP methods.

Ultimately, the proposed DP method trades a mathematically negligible single-step approximation error for a computational speedup of up to three orders of magnitude, all while maintaining the exact same long-term financial performance as the traditional exact MILP approach.

4.3. Yearly high-frequency rolling intrinsic trading

Leveraging the speed of our proposed DP method, we present full-year DP-RI trading results under a standard battery and simulation setup for the year 2021 in Fig. 4, and compare them against two materially different, implementable industry benchmarks: The first is a static Day-Ahead (DA) benchmark, which optimizes the battery schedule once per day based on the cleared DA auction prices. The second is an implementable continuous intraday benchmark based on Volume-Weighted Average Prices (VWAPs). Adapted from a state-of-the-art industry battery index [5], this strategy calculates VWAPs from transaction data in trailing 15-minute intervals and optimizes the schedule based on these aggregated price curves at that frequency. To ensure a fair and consistent multi-day evaluation, we adapted the published index by

Table 3
Performance comparison for four representative weeks in 2021 between the rolling intrinsic solved using the benchmark MILP and three DP approaches with different approximations due to their value function granularities. All results are rounded to significance.

		Reward [€]	Sim-Time [h]	Cycles/Day	Solves	Traded Vol. [MWh]
Feb 01–14	MILP	7895	6.77	2.3	959,900	1901
	DP-101	7538	0.43	2.3	959,900	1825
	DP-51	8051	0.22	2.2	959,800	1905
	DP-11	7851	0.05	2.2	959,500	1878
Apr 01–14	MILP	9487	15.81	2.5	733,000	1856
	DP-101	9434	0.34	2.5	733,100	1817
	DP-51	9353	0.17	2.5	733,000	1792
	DP-11	9543	0.04	2.5	733,200	1872
Jul 01–14	MILP	11,565	17.07	2	901,000	1948
	DP-101	11,161	0.43	2	901,000	1873
	DP-51	11,158	0.22	2	900,900	1884
	DP-11	11,181	0.05	2	901,000	1899
Oct 01–14	MILP	28,553	56.93	2.6	1,012,900	3350
	DP-101	28,134	0.47	2.6	1,012,500	3231
	DP-51	27,952	0.24	2.5	1,012,600	3190
	DP-11	28,014	0.06	2.5	1,012,600	3236

optimizing across all actively tradable products without artificial end-of-day state-of-charge constraints, integrating our linear degradation costs, and utilizing an MILP solver to correctly handle negative prices and enforce minimum bid-size increments. To make this strategy implementable on a continuous market, we submit the resulting target positions as market orders to the actual historical limit order book at the time of solution (starting trading at 11 pm for the next day's delivery to avoid the extreme illiquidity present at the 3 pm gate opening).

Fig. 4 illustrates the cumulative rewards of these benchmarks against our proposed LOB-based strategy. When the VWAP strategy is forced to interact with the actual LOB, rather than assuming infinite liquidity at the theoretical average price, it systematically overestimates available volumes and suffers from significant price slippage upon execution. While the LOB strategy re-optimizing every 15min yields profits roughly comparable to the static DA benchmark, our high-frequency LOB strategies demonstrate a clear performance advantage. By accurately modeling depth-dependent pricing and capturing short-lived opportunities through millisecond-level reoptimization, our fastest LOB strategy outperforms the realistically implemented VWAP benchmark by more than 36% in realized annual profits, highlighting the critical value of detailed microstructural modeling of the continuous intraday market.

The DP settings used for the comparisons are outlined in Section 4.1 with a value function discretization setting of $m = 11$ and a technical delay of 200 ms plus measured actual solve times. We evaluate performance starting from the baseline with a solve-frequency of 60 minutes, i.e., solving the intrinsic optimization once every hour, which results in a reward of €221k (8.7k orders submitted at the exchange) and a simulation runtime of 5.6 minutes, compared to the strategy that solves it with every relevant order book update (apart from the waiting durations introduced previously) yielding a reward of €349k (30k orders submitted at the exchange) and a simulation runtime of 86 minutes. We define any order message submitted to the LOB as relevant, if it is placed at the head of its corresponding bid/ask stack. These results clearly demonstrate that solving the intrinsic at higher frequencies leads to progressively increasing profits as the trading frequency rises. In this case, a 58% increase between our slowest and quickest solve frequency. We observe a clear increase in the slope of the cumulative reward starting in October 2021. This most likely originates from higher and more volatile gas prices in Europe during this time, which subsequently translates to more extreme intraday electricity prices, increasing the revenue potential of storage.

To extend the information presented in Fig. 4's inset, Fig. 5 depicts the detailed sequential decisions made by our policy on the last day of the inset. Looking closely at the overlapping dots, one can observe that some trading decisions are only a few hundred milliseconds apart.

An analysis of the daily profit distribution reveals a highly favorable, asymmetric risk profile inherent to the myopic rolling intrinsic policy. The algorithm effectively capitalizes on extreme scenarios characterized

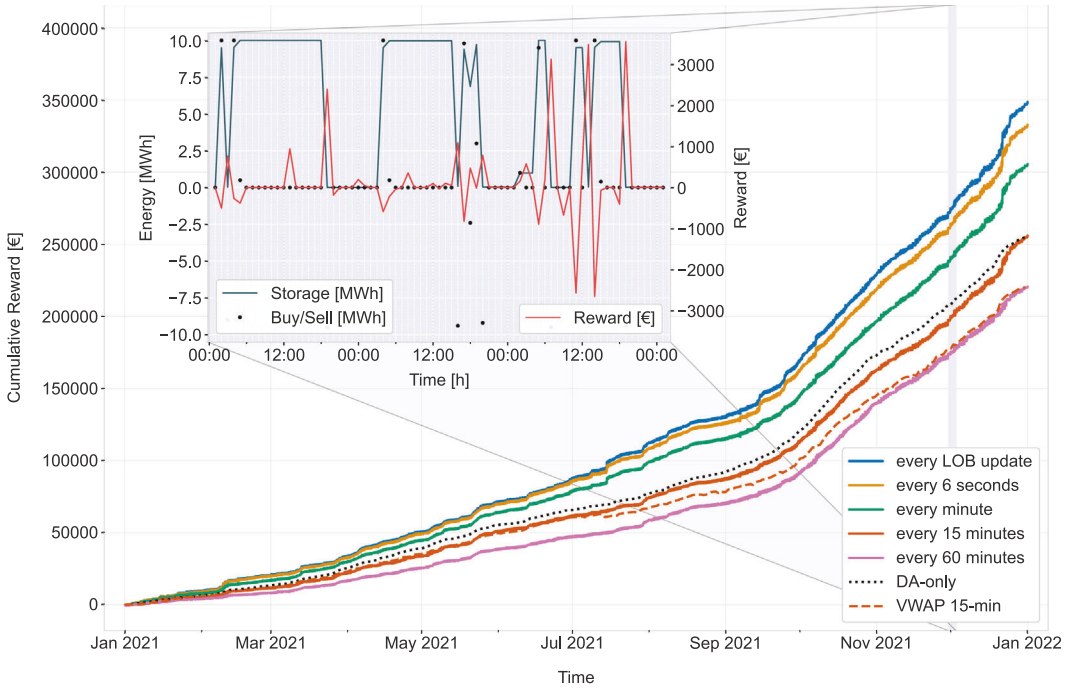


Fig. 4. The cumulative reward of the intrinsic rolling over the entire year 2021 for various intrinsic-solve frequencies, compared against two benchmark strategies. The inset shows the battery operation over the span of three arbitrarily selected days for the strategy that solves the intrinsic at every relevant LOB update (taking into account technical delays), with the final schedule determined by the RI.

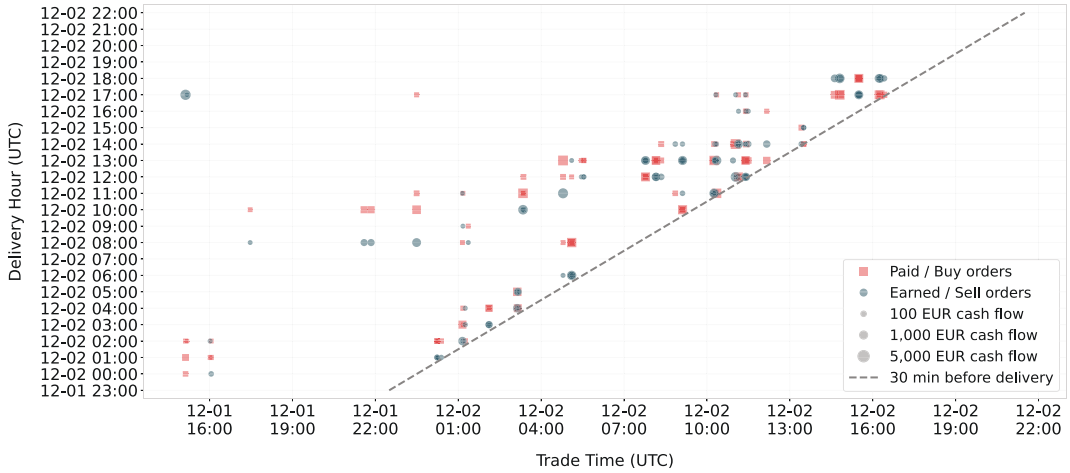


Fig. 5. Detailed trading decisions for products traded on 2021-12-02 (Berlin time), which is the last day depicted in Fig. 4's inset.

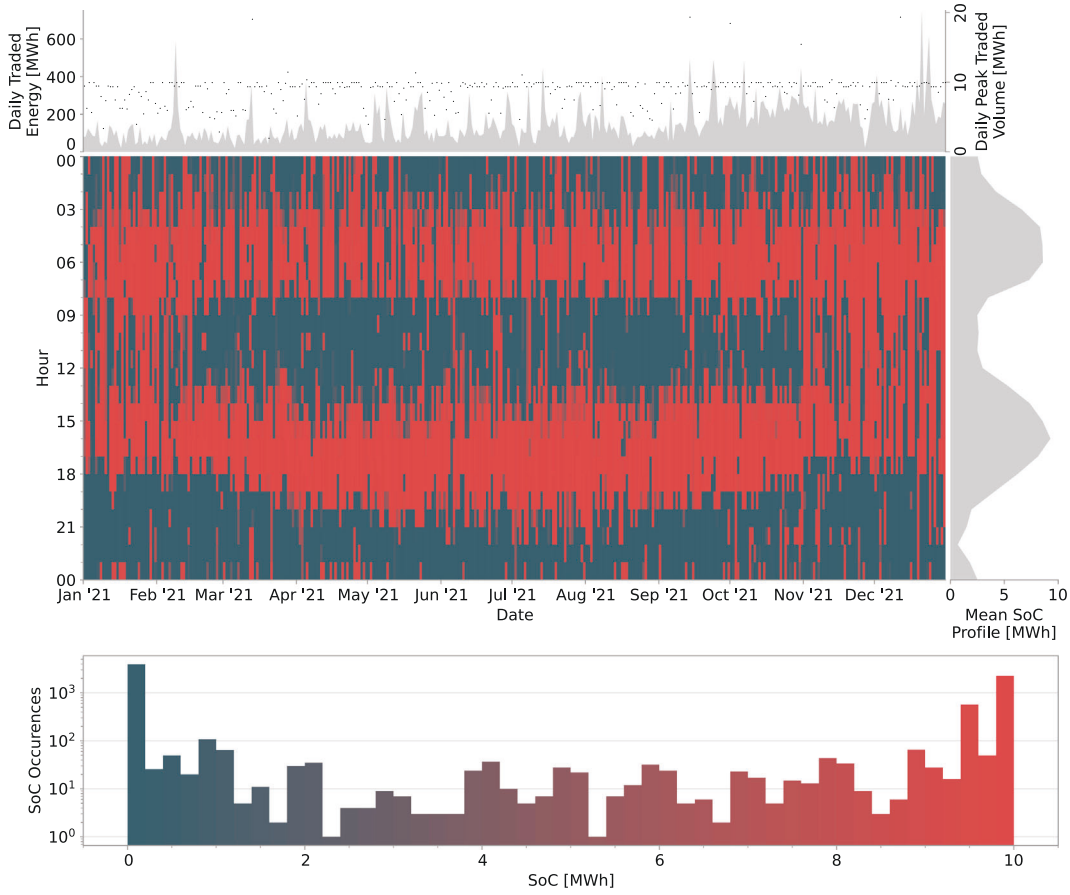


Fig. 6. Yearly heatmap of the final battery schedule for the year 2021 determined by the RI. The upper plot shows the final state of charge (SoC) of the battery for each hour of the year, with the color gradient defined in the lower plot. The gray areas show the total traded energy per day and the yearly mean SoC profile (average SoC for each hour 1–24 of the day across the entire year), while the dots in the upper inset represent the largest trade per day submitted by the RI. We observe a clear preference for charging the battery in the early morning and afternoon hours. The lower plot shows the logarithmic distribution of SoC states of the battery, as determined by the RI's schedule. Figure adapted from [6].

by high price volatility, resulting in a heavy right tail of highly profitable days (most notably in late 2021). Conversely, the strategy is virtually immune to extreme daily losses. During periods of low liquidity or unfavorable prices, the policy simply refrains from trading rather than speculating.

Figs. 6 and 7 illustrate DP-RI decisions, using the same settings and our quickest strategy, solving the intrinsic at every relevant LOB update. Fig. 6 provides an overview of the annual battery state-of-charge (SoC) schedule shaped by the RI's trading decisions, along with the daily trading volumes. We observe a distinct average SoC pattern, suggesting a

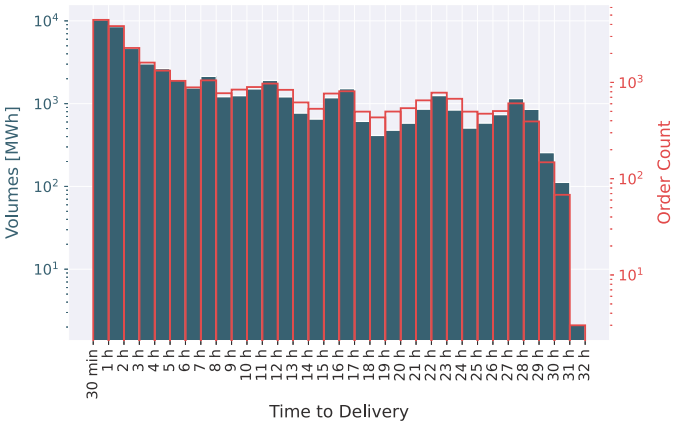


Fig. 7. Logarithmic order submission-time analysis of the RI simulated over the entire year 2021.

Table 4
Statistics on missed orders and subsequent feasibility corrections over the year 2021. 'Costs' represent the foregone profit impact of correcting non-physical positions caused by execution latency.

	Corrections				Time (sec)	
Type	Count	% of solves	Costs (€) / % of total	Sim Reward (€)	Mean	Median
DP-11	867	7.3%	−25,200 / −7.2%	349,000	0.93	0.30
DP-51	996	7.1%	−21,400 / −6.1%	351,000	8.61	0.30
DP-101	989	6.9%	−24,000 / −6.8%	352,000	1.96	0.30

strong daily periodicity of arbitrage opportunities, with higher charge levels during the early morning and afternoon hours. This can be seen by looking at the trading decisions in the central heatmap of trading decisions as well as the mean SoC profile.

The lower plot shows the frequency of different SoC levels throughout the year. Given that our storage has a duration of one hour, one would expect the storage to be either completely full or completely empty most of the time. However, the SoC occupies levels in between these extremes surprisingly often. This is most likely caused mainly by complex trading behavior induced by nonlinearities in order book-based trading, which makes simple bang-bang strategies suboptimal. This highlights the difference between a naive price-taking implementation of the rolling intrinsic and our more realistic approach.

Lastly, when looking at the daily traded energy over the year on the top of the figure, we see increased trading activity towards the end of the year, which fits the increased slope of cumulative profits for this time period observed in Fig. 4.

Fig. 7 shows the number of orders executed by the RI, classified by their proximity to physical delivery. Comparing this to Fig. 2 which shows the number of orders placed in the market, we see that the pattern is much more uniform. This is likely due to the fact that the rolling intrinsic tries to capitalize on time spreads which often induces trading pairs of products with large differences in the time to delivery, e.g., when buying cheap in the early morning to sell expensive in the evening. This induces the RI to frequently trade products that are not yet liquid, which is one of the biggest weaknesses caused by the myopic nature of the policy.

Finally, as discussed in Section 4.1, market latency can lead to execution failures for our submitted limit orders. When an order is missed because the order book has already shifted during the combined solve time and technical delay, the algorithm may temporarily record a non-physical position that it must correct in the immediately following optimization step. Table 4 details the frequency, time-to-resolution, and financial impact of these feasibility corrections over the full-year simulation. Across different dynamic programming discretizations, approximately 7% of the orders were missed due to latency. The rolling policy naturally resolved these infeasibilities rapidly, with a median resolution time of 0.3 s. However, these missed opportunities and subsequent corrections carry a tangible cost, reducing the total theoretical simulation profits by roughly 6% to 7%. The consistency of these statistics across different state grid sizes ($m = |G|$) aligns with expectations, as the measured solve time for all discretizations remains far below the fixed 200 ms technical delay, making the latter the main reason for missing orders. Isolating the exact financial penalty of a missed order is non-trivial, as subsequent re-optimizations might replace the order, abandon it, or combine it with new trades, sometimes even at more favorable prices. To estimate the foregone profit impact reported in Table 4, we employ a volume-proportional scaling approximation. If an order fails to match, we examine the replacement orders placed in the immediate next solve that share the same trade direction. We then scale the cost comparison by the volume share of the missed order. For example, if the algorithm misses a 1.0 MWh buy order at 10 €/MWh, and the subsequent solve proposes a 1.3 MWh buy order at 11 €/MWh, we estimate the correction loss based on the original 1.0 MWh volume: $((1.3 \times 11)\frac{1.0}{1.3} - 1.0 \times 10) = 1$ €.

This method provides a realistic approximation of the friction introduced by execution latency without attempting to track ambiguous cascading correction-intents across multiple future solves. We acknowledge, however, that this volume-proportional scaling is a heuristic that introduces a degree of uncertainty. Due to the highly path-dependent nature of the rolling intrinsic policy, a missed order does not simply force a 1-to-1 replacement trade. It can occasionally cause the algorithm to abandon a trade entirely or shift its strategy to different delivery hours. Consequently, the true opportunity cost of a missed order could be slightly higher or lower than our estimate. Nevertheless, because this approximated latency penalty, which accounts for a roughly 6% - 7% reduction in total theoretical profit, is already fully incorporated into our reported trading results, this uncertainty does not threaten the robustness of our core conclusions. Even when absorbing these estimated latency frictions, our high-frequency DP approach continues to significantly outperform the lower-frequency and day-ahead benchmarks.

4.4. Parameter evaluation

The DP rolling intrinsic optimization relies on numerous parameters, each influencing trading behavior and solution accuracy in different ways. This section compares key storage and optimization parameters, demonstrating that our year-long trading results remain robust and that the parameters' effects align with expectations.

To achieve this, we compare the effects of varying injection-efficiencies (assuming $\eta^+ = \eta^-$), linear degradation costs ν_{deg} and maximum storage capacities $\bar{s}$, each for three different equidistant DP value function grids G with granularity $|G| = m$. Table 5 provides a shortened overview of yearly rewards for each of the optimizations. Two clear trends can be observed as expected: A finer grid G (larger m) results in more profitable trading as the solution is more accurate, and more penalizing injection efficiencies η and linear degradation costs ν result in lower rewards because fewer trading opportunities become profitable. Both parameters have a significant effect on profits, with the most favorable settings nearly doubling the profits of the least favorable setting. This is especially interesting for the case of degradation costs, which are likely to decrease significantly in the near future, thereby greatly increasing the profitability of short-term trading on intraday markets.

Additionally, batteries with a smaller maximum storage capacity produce larger profits per MWh storage capacity, which is mainly caused by limited market depth and liquidity. Operating a larger storage system requires executing larger trade volumes. To fill these larger volumes, the algorithm must sweep deeper into the limit order book, consequently

Table 5
A reward comparison of simulation results for the full year 2021, using four different storage parameter settings. All simulation base-settings are equal to the setting defined in Section 4.1, setting the ping-delay to 200 ms and using actual measured solve-times at execution. Varying storage parameters has the expected effect on yearly rewards, where a finer DP storage discretization generally leads to marginally higher rewards, increasing degradation costs and losses reduces the rewards, and operating a larger storage yields lower rewards per storage capacity. All results are rounded to significance, and the specific battery reward on the right table is given in units of €/MWh.

	$x\eta^+ = \eta^-$					$\bar{s}$ [MWh] 1h-battery			
Reward [€]	0.9	0.95	0.99	1.0	Reward/$\bar{s}$ [$\frac{€}{MWh}$]	5	10	20	40
DP-11	277,000	349,000	441,000	478,000	**DP-11**	36,600	34,900	32,600	29,700
DP-51	278,000	351,000	444,000	482,000	**DP-51**	36,900	35,100	33,200	30,400
DP-101	279,000	352,000	442,000	483,000	**DP-101**	36,900	35,200	33,400	30,600
	ν_{deg} [€/MWh]					$\bar{s}$ [MWh] 2h-battery			
Reward [€]	8	4	2	0	Reward/$\bar{s}$ [$\frac{€}{MWh}$]	5	10	20	40
DP-11	263,000	349,000	416,000	434,000	**DP-11**	27,300	27,000	25,700	23,900
DP-51	264,000	351,000	419,000	439,000	**DP-51**	27,100	27,400	26,200	24,600
DP-101	265,000	352,000	419,000	441,000	**DP-101**	27,700	27,300	26,400	24,700

accepting progressively worse prices and shrinking the captured price spreads. Similarly, a battery with a duration of 2h yields lower profits, but more than half of the profits observed for our standard setup with a duration of 1h, because it can no longer charge and discharge in the hour(s) of the singular best price(s) of the day.

4.5. Parametrization of the rolling intrinsic

The rolling intrinsic strategy is suboptimal due to its myopic nature. In particular, the lack of foresight leads to *impatient trading decisions* that use flexibility early on, thereby foregoing parts of the storage's revenue potential. This happens in particular when the RI trades products with limited liquidity, typically long before gate closure, and for rather disadvantageous prices.

This opens the door for parametric variations of the rolling intrinsic policy that aim to improve on these issues by adjusting the trading behavior. However, finding optimal parameters requires extensive tuning efforts and, therefore, only works in combination with an efficient implementation of the policy. Our method therefore opens the door to optimized variations of the rolling intrinsic, which can be trained within a realistic timeframe.

To illustrate this point, we introduce an empirically motivated unitless linear parameter ϕ to the intrinsic optimization. It acts as a penalty for products with a large bid-ask spread. We then adapt the optimization (6) by extending π_t to $\hat{\pi}_t = \pi_t - (\phi \cdot \delta_{t^*,t}) \mid f_t^0 - f_t \mid$, with $\delta_{t^*,t}$ defined as the bid-ask spread of product t at time t^*. Intuitively, it becomes clear that ϕ incentivizes the rolling intrinsic to trade at times of higher market liquidity, i.e., at times closer to delivery, to reduce the chance of trading less profitably too early.

The increased speed of our DP solution method allows us to quickly search for optimal parameters ϕ. For the sake of simplicity and just to show a proof of concept, we fix the parameters ϕ at constant values for each trading month, and use the previous month $m-1$ to train ϕ for a given month m. When evaluating this strategy over the entire year 2021, this sliding window parameter training results in 12 separate training windows. We train the parameters on the highest intrinsic-solve frequency (see Section 4.3) using DP and simulation settings outlined in Section 4.1, which represents the most realistic training scenario. We treat our monthly simulation rewards as a black box function $f(\phi)$ and, given its simple dependence on a singular parameter, use Brent's method to find its maximum.

The resulting optimal parameters for the year 2021 have an average penalty of $\bar{\phi} = 2.1 \pm 1.6$. Comparing year-long out-of-sample results of the newly trained policy yields an increase of 8.4% in trading reward over the entire year 2021, accumulating €376k, compared to the €347k reward of the standard policy with the penalty set to 0, i.e., $\phi \equiv 0$.

While this specific extension remains an illustrative tuning exercise rather than a comprehensive new trading framework, it highlights the immense practical value of the proposed fast DP approximation. Ultimately, it demonstrates that our method provides the computational foundation required to rapidly prototype, iteratively train, and backtest more sophisticated parametric or machine-learning-driven strategies over massive, high-frequency datasets, a task that would be strictly impossible using traditional formulations.

5. Conclusion

In this paper, we present a novel method to compute the rolling intrinsic policy for trading with a battery on continuous intraday markets. Our method achieves a solution speed improvement of up to three orders of magnitude while maintaining a sufficient level of accuracy. Our work lays the foundation for extensive future research through the extension and adaptation of our rolling intrinsic formulation and implementation, both for researchers and industry.

Our findings demonstrate the importance of strategies that take into account the exact workings of the continuous intraday market and, at the same time, can be executed fast enough to keep up with the millisecond pace of the market. We show through realistic backtesting that a solution frequency of at least every second increases rolling intrinsic profits substantially compared to solving the intrinsic only every few minutes or even every hour. Our synthesis of the backtesting results underscores a fundamental market paradigm: the continuous clearing mechanism structurally rewards high-frequency execution. Strategies that fail to react to tick-level order book updates frequently leave value on the table, as delayed optimization limits a battery's ability to lock in short-lived price spreads.

Furthermore, our efficient implementation enables us to backtest the profitability of the rolling intrinsic policy over extended time periods, which requires a substantial number of intrinsic solves across the simulation period, making traditional MILP or LP methods impractical due to their comparatively slow solution times. This also makes it possible to train parametric extensions of the rolling intrinsic that correct some of its shortcomings, as we demonstrate by optimizing a penalty for trading illiquid products, which substantially increase trading profits.

For practitioners, our results illustrate that resolving the limit order book at high frequencies is a highly relevant factor when designing and backtesting intraday trading algorithms. Our empirical analysis focuses exclusively on the 2021 German continuous intraday market, which operates within the European Single Intraday Coupling. While absolute profitability will naturally vary across different historical years, market regimes, and specific international regulatory environments, our findings clearly highlight the structural mechanics of continuous clearing. Within this specific market design, strategies equipped with the computational speed to process and react to tick-level data possess a distinct execution advantage over slower, lower-frequency benchmarks.

We acknowledge specific limitations in our modeling approach, particularly the linear approximation of battery degradation and the discretization error inherent in the DP formulation relative to the exact MILP. Furthermore, while our use of full-depth order book data ensures high fidelity, the backtest remains a simulation. It inherently cannot capture endogenous market impact, specifically how other market participants might have dynamically adjusted their strategies in response to our orders. Lastly, a notable limitation of our current approach is the exclusive focus on the intraday market, which omits the co-optimization potential with other highly profitable platforms such as the day-ahead or frequency containment reserve markets.

An important future extension of our work lies in more sophisticated parametrizations of the rolling intrinsic. Our naive method could easily be adapted to multiple additional parameters like trading delays, affecting various aspects of the RI's trading behavior, further utilizing the speed of our optimization method. Additionally, incorporating nuanced degradation costs, such as State of Health (SoH) and non-linear depth-of-discharge, would significantly enhance the realism and applicability for practitioners. Furthermore, future work could naturally extend the trading to a multi-market scenario, where additional decision variables expand the scope of optimization to multiple subsequent electricity markets, e.g., the reserve or day-ahead markets.

Finally, extending our intrinsic optimization with dynamic price forecasts, or more general value functions of future actions, would greatly improve the RI's profitability, as the optimization would be able to anticipate future price changes and capture the option value of waiting, rather than relying on strictly myopic execution. Fine-tuning non-myopic, forecast-driven, or stochastic models to operate at the high frequencies enabled by our DP approach is a logical and highly promising avenue for future research.

CRediT authorship contribution statement

David Schaurecker: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **David Wozabal:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Data curation, Conceptualization. **Nils Löhndorf:** Software, Methodology, Formal analysis, Conceptualization. **Thorsten Staake:** Writing – review & editing, Project administration, Conceptualization.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used Gemini and ChatGPT to improve language and readability. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors thank Markus Kreft and Patrick Langer for the helpful methodological discussions and for their support with some of the figures. In addition, we thank Simon Hirsch for the factual checks, general feedback and support in wrapping the package to Python and setting up the documentation.

Appendix A. Code publication

We publish our method as a Python package *BitePy* (Battery Intraday Trading Engine), where users can preprocess their intraday market data, set battery and DP parameters, run simulations and analyze results. It is hosted on PyPi and can be easily installed via *pip install bitepy*. Detailed documentation and tutorials on the package can be found on GitHub.

Data availability

We are not allowed to share market data, but appended a link to our simulation engine's GitHub at the end of the manuscript.

References

[1] Aid R, Gruet P, Pham H. An optimal trading problem in intraday electricity markets. Math Financ Econ 2016;10:49–85.
[2] Bertrand G, Papavasiliou A. Adaptive trading in continuous intraday electricity markets for a storage unit. IEEE Trans Power Syst 2020;35:2339–50. https://doi.org/10.1109/TPWRS.2019.2957246
[3] Boukas I, Ernst D, Théate T, Bolland A, Huynen A, Buchwald M, Wynants C, Cornélusse B. A deep reinforcement learning framework for continuous intraday market bidding. Mach Learn 2021;110:2335–87. https://doi.org/10.1007/s10994-021-06020-8
[4] Braeuer F, Rominger J, McKenna R, Fichtner W. Battery storage systems: an economic model-based analysis of parallel revenue streams and general implications for industry. Appl Energy 2019;239:1424–40. https://doi.org/10.1016/j.apenergy.2019.01.050
[5] Brucksch J, van Ouwerkerk J, Sauer DU. ISEA battery revenue index; 2025 https://battery-charts.de/revenue-index/ [Accessed: 27 May 2026].
[6] Brudermueller T, Kreft M. Smart meter data analytics: practical use-cases and best practices of machine learning applications for energy data in the residential sector. In: ICLR 2023 workshop on tackling climate change with machine learning; 2023. https://www.climatechange.ai/papers/iclr2023/3.
[7] Cole W, Karmakar A. Cost projections for utility-scale battery storage: 2023 update. Technical Report NREL/TP-6A40-85332 Golden, CO: National Renewable Energy Laboratory; 2023. https://www.nrel.gov/docs/fy23osti/85332.pdf.
[8] Cornejo M, Graner M, Hesse H, Jossen A. Evaluating the impact of model accuracy for optimizing battery energy storage systems. arXiv:2506.17059. 2025.
[9] EPEX SPOT SE. EPEX SPOT annual reports; 2017-2023 https://www.epexspot.com/ [Accessed: 17 June 2024].
[10] Epex Spot SE. Press release: intraday auctions (IDAs) were implemented across Europe on 13 June 2024; 2024. https://www.epexspot.com/en/news/intraday-auctions-idas-were-implemented-across-europe-13-june-2024.
[11] Faria E, Fleten SE. Day-ahead market bidding for a Nordic hydropower producer: taking the elbas market into account. Comput Manag Sci 2011;8:75–101.
[12] Fleten SE, Kristoffersen TK. Short-term hydropower production planning by stochastic programming. Comput Oper Res 2008;35:2656–71.
[13] Goldman Sachs. Electric vehicle battery prices are expected to fall almost 50% by 2026; 2023 https://www.goldmansachs.com/insights/articles/electric-vehicle-battery-prices-are-expected-to-fall-almost-50-percent-by-2025 [Accessed: 21 December 2024].
[14] Graf C, Kuppelwieser T, Wozabal D. Frequent auctions for intraday electricity markets. Energy J 2024;45:231–56. https://doi.org/10.5547/01956574.45.1.cgra
[15] Gray J, Khandelwal P. Realistic gas storge models II: trading strategies. Commod Now Sep 2004;1–5.
[16] Gray J, Khandelwal P. Towards a realistic gas storage model. Commod Now 2004;7:1–4.
[17] Gräf D, Marschewski J, Ibing L, Huckebrink D, Fiebrandt M, Hanau G, Bertsch V. What drives capacity degradation in utility-scale battery energy storage systems? The impact of operating strategy and temperature in different grid applications. J Energy Storage 2022;47:103533. https://doi.org/10.1016/j.est.2021.103533
[18] Gurobi Optimization, LLC. Gurobi optimizer reference manual; 2024.
[19] IRENA. Electricity storage and renewables: costs and markets to 2030. Accessed: 2017. https://www.irena.org/publications/2017/Oct/Electricity-storage-and-renewables-costs-and-markets [16 December 2024].
[20] Jiang D, Powell W. Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. INFORMS J Comput 2015;27:525–43.
[21] Koch C, Hirth L. Short-term electricity trading for system balancing: an empirical analysis of the role of intraday trading in balancing Germany's electricity system. Renew Sustain Energy Rev 2019;113:109275. https://doi.org/10.1016/j.rser.2019.109275
[22] Kraft E, Russo M, Keles D, Bertsch V. Stochastic optimization of trading strategies in sequential electricity markets. Eur J Oper Res 2023;308:400–21. https://doi.org/10.1016/j.ejor.2022.10.040
[23] Kuppelwieser T, Wozabal D. Intraday power trading: toward an arms race in weather forecasting? OR Spectr 2023;45:57–83.
[24] Lai G, Margot F, Secomandi N. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. Oper Res 2010;58:564–82.
[25] Lai G, Wang M, Kekre S, Scheller-Wolf A, Secomandi N. Valuation of storage at a liquefied natural gas terminal. Oper Res 2011;59:603–16.
[26] Löhndorf N, Wozabal D. Gas storage valuation in incomplete markets. Eur J Oper Res 2021;288:318–30.
[27] Löhndorf N, Wozabal D, Minner S. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. Oper Res 2013;61:810–23.
[28] Löhndorf N, Wozabal D. The value of coordination in multimarket bidding of grid energy storage. Oper Res 2023;71:1–22. https://doi.org/10.1287/opre.2021.2247

[29] Martínez-Barbeito M, Gomila D, Colet P. Dynamical model for power grid frequency fluctuations: application to Islands with high penetration of wind generation. IEEE Trans Sustain Energy 2023;14:1436–45. https://doi.org/10.1109/TSTE.2022.3231975
[30] Neuhoff K, Ritter N, Salah-Abou-El-Enien A, Vassilopoulos P. Intraday markets for power: discretizing the continuous trading? Discussion Paper 1544 DIW Berlin; 2016. https://doi.org/10.2139/ssrn.2723902. https://ssrn.com/abstract=2723902.
[31] Regelleistung Online. BESS revenue index 1H; 2024 https://www.regelleistung-online.de/german-energy-storage-revenue-index/bess-revenue-index-1h/ [Accessed: 22 December 2024].
[32] Sage M, Campbell J, Zhao YF. Enhancing battery storage energy arbitrage with deep reinforcement learning and time-series forecasting. In: ASME 2024 18th international conference on energy sustainability. American Society of Mechanical Engineers; 2024. https://doi.org/10.1115/ES2024-130538
[33] Saldarini A, Longo M, Brenna M, Zaninelli D. Battery electric storage systems: advances, challenges, and market trends. Energies 2023;16.
[34] Seifert PE, Kraft E, Bakker S, Fleten SE. Coordinated trading strategies for battery storage in reserve and spot markets. arXiv:2406.08390. 2024.
[35] Semmelmann L, Dresselhaus J, Miskiw KK, Ludwig J, Weinhardt C. An algorithm for modelling rolling intrinsic battery trading on the continuous intraday market. SIGENERGY Energy Inf Rev 2025;4:163–74. https://doi.org/10.1145/3717413.3717428
[36] Ullah F, Zhang X, Khan M, Mastoi MS, Munir HM, Flah A, Said Y. A comprehensive review of wind power integration and energy storage technologies for modern grid frequency regulation. Heliyon 2024;10:e30466. https://doi.org/10.1016/j.heliyon.2024.e30466
[37] Wozabal D, Rameseder G. Optimal bidding of a virtual power plant on the Spanish day-ahead and intraday market for electricity. Eur J Oper Res 2020;280:639–55.
[38] Xu B, Zhao J, Zheng T, Litvinov E, Kirschen DS. Factoring the cycle aging cost of batteries participating in electricity markets. arXiv:1707.04567. 2017.
[39] Zhang Y, Ridinger W, Wozabal D. Joint bidding on intraday and frequency containment reserve markets. arXiv:2510.03209. 2025.