BENV0096: MSc ESDA Dissertation

"Evaluating multi-market wholesale arbitrage bidding strategies of battery energy storage portfolios of different compositions."

by

Adam Lakatos

14th February, 2025

Paper submitted in part fulfilment of the
Degree of Master of Energy Systems and Data Analytics

Energy Institute
University College London

Word count: 9,979

I want to thank my supervisor, Nikolaos Chrysanthopoulos, for his kindness and support, he had a transformative impact on the process of writing this dissertation. I also want to thank Modo Energy for providing access to their Modo Terminal, helping me get to grips with the mechanics of electricity- and energy storage markets.

Table of Contents

Nomenclature

$q_{m,t,c}$	Charge price on market m, in period t, after accounting for degradation costs, in £/MWh.
$q_{m,t,d}$	Discharge price on market m, in period t, after accounting for degradation costs, in £/MWh.
$p_{m,t}$	Clearing price on market m, for period t, in £/MWh.
x_{BESS}	Constant degradation cost per MWh.
$c_{m,t}$	Contracted amount of charge on market m, for period t, in MWh.
$d_{m,t}$	Contracted amount of discharge on market m, for period t, in MWh.
E_{max}	Maximum energy delivery capacity of storage asset within a single trading period, in MWh.
SoC_t	State of Charge of storage, in MWh, after physical charging and discharging actions in period t.
SoC_{limit}	Total energy storing capacity of storage, in MWh.
$cycle_t$	Amount of total charge delivered up to and including period t, in MWh.
$cycle_{limit}$	Amount of total energy uptake allowed per day by the daily cycling limit, in MWh.
P_{BESS}	Maximum power delivery capacity of battery storage system, in MW.
D_{BESS}	Duration of the battery: maximum amount of time the battery can charge for at maximum power, in hours.
$\#_{cycles}$	Daily limit of cycles, where one cycle is a full charge and discharge of the battery.
net_t	Final physical delivery requirements of storage in period t after all market positions have been netted, in MWh.
ac_t	Auxiliary variable equal to net charge amount if net charge is positive, 0 otherwise, in MWh.
δ_t	Auxiliary dummy variable.
$c_{stage1,t}$	Contracted charging volume in period t, from first stage DA bidding, in MWh.
$d_{stage1,t}$	Contracted discharging volume in period t, from first stage DA bidding, in MWh.
T	Number of time periods in optimisation.

1. Introduction

The ongoing increasing penetration of variable and non-dispatchable renewable energy sources – driven by the global effort to decarbonise power grids – increased the need for fast response services that can help in correcting the imbalances caused by the scheduling inaccuracies of such renewables. Solar PV power in particular has a distinct daily generation pattern, that further increases daily spreads and the occurrence of extreme price events given abundant cheap solar power during daytime and the need to dispatch more expensive generators in the evening times. As the levelized cost of storage solutions declined, they have emerged as powerful solutions to meet the market's needs for rebalancing. Storages have a wide range of variability when it comes to their scale, efficiency or response time, and possess unique capabilities like bidirectional power flow, cost-free stand-by operation, and rapid response, which result in a wide versatility of use-cases and revenue sources. In Great Britain (GB), grid-connected utility-scale energy storage systems can effectively participate in frequency response services to maintain grid stability by providing energy on short notice when the frequency on the grid increases due to excess demand, or by charging from the grid when there is too much generation. Behind-the-meter residential storages are more ideal for demand shifting price response or backup generation in co-location with a renewable generation source, while energy storage price arbitrage opportunities on Wholesale- or Balancing Markets (BM) incentivises market participants to charge during periods of low prices and to discharge during periods of high prices, earning profits on the margins. Arbitraging stabilises the system by bringing prices across time and across markets closer together on a given day. Batteries can earn additional revenues from participating in Reserve and Capacity Markets, which are longer horizon than frequency response markets and pay batteries for guaranteed availability in case activation is required. Usually, for batteries to be economically viable, multiple of these applications are exploited concurrently, sometimes selling the same capacity across multiple markets, called "value stacking".

The Battery Energy Storage (BES) market grew in GB from a total installed capacity of 1.2GW in Q4, 2021 to 4.7GW in Q4, 2024, in the context of a total of 105GW of installed capacity from all types of generation sources (Department for Energy Security and Net Zero, 2024). As a result, although 84% of an average battery's revenue stack in 2022 in GB was from providing frequency response services, by 2024 it was only 20% since frequency response markets are limited in size and the rapid buildout of new storage systems quickly saturated the markets, driving prices down to one seventh of what they were in 2022. As a result, current market outlooks forecast that 86% of future revenues of a newly built battery system will stem from Wholesale and BM energy trading (Modo Energy, 2024a).

In Great Britain, Wholesale markets run from more than 24 hours up to 15 minutes before the time of delivery. The most significant are the Day-Ahead (DA) auction markets which simultaneously clear a uniform price for all auction participants for each hour, or half-hour of the day, before start of day. The DA markets have a combined average daily traded volume of 500GWh compared to the 60GWh contracted on the Intra-Day (ID) market (Modo Energy, 2023b). Its clearing process is determined by the demand requirements of retail suppliers – driven by expected behavioural and environmental conditions – and the operational costs and capabilities of generators based on gas, carbon, or coal prices, solar and wind conditions or interconnector activity. In contrast, the ID market depends more on unexpected events within-day such as generator unavailability, or weather changes and is used speculatively to hedge against System Prices. It relies on an order-book architecture that continuously matches bids and offers up to 15 minutes before delivery, where each individual trade can have a different clearing price, but any trade requires a party willing to take the opposite side. Overall, as a consequence of the higher liquiditiy, DA prices tend to have a smoother "duck-curve" shape, while the ID prices can diverge significantly and tend to have higher spreads. Meanwhile, on the BM – which also clears a day ahead and is responsible for balancing the market in real-time by acquiring and activating reserves – spreads between the Bid and Offer prices were 1.8 times higher than spreads on the DA markets in 2023 (Modo Energy, 2024b). However, unlike wholesale markets which are self-dispatched, BM resources are dispatched by the Transmission System Operator (TSO) and payout is not guaranteed, introducing added uncertainty. Storage operators focusing on arbitrage trading have an opportunity to coordinate their bidding activity across these markets, earning on price spreads and adjusting positions while taking into account the sequentiality of market clearings and the deviations in market dynamics. Optimal trading strategies in such a multi-stage decision-making problem with time coupling constraints requires making sequential capacity commitments to maximise profits while forecasting and hedging against uncertainties in future stages. Asset-backed strategies relying on storage systems are able to benefit from non-physical trading of intermarket spreads – when multiple prices exist for the same delivery contract – while still having the option to cut the trade short and simply deliver the first position.

This paper develops single market and multi-market arbitrage trading strategies for the DA and ID markets with different degrees of coordination, under price uncertainty, with the objective of evaluating i) the potential returns achievable by the different strategies, ii) their relative adaptability to changing market conditions, iii) the degree to which online formulations are able to capture potential profits iv) and the effect of modifying the physical configurations of the battery. Performance is benchmarked by generating results for fully optimised offline formulations, while multiple variations of price uncertainty and forecasting methods are explored within online formulations in order to quantify each component's impact on the results of the bidding process. The forecasting models are trained, and the strategies are backtested on historical data from the GB electricity markets. The result is an online strategy that successfully integrates the biddings across markets. The remainder of the paper is divided into several sections: Section 2 reviews the literature, Section 3 details the methodology for developing the bidding strategies, Section 4 describes the data, Section 5 presents and evaluates the results of the optimisations and Section 6 concludes.

2. Literature Review

There is a well-established body of literature on the topic of arbitrage trading backed by a battery energy storage system, though there are significant differences in the modelling methodologies used across papers, largely due to the diversity of market and operational framework designs that papers can set out in their experiments. Differences can stem from only focusing on one of the Wholesale markets, or optimising their combination; the form of degradation cost function considered; whether price uncertainty is incorporated; whether DA bidding is quantity-only or price-quantity; whether the ID market is treated as operating on a sequential Real-Time or a continuous order book framework; whether there are co-located renewables or additional markets, such as frequency response; or whether the objective is live implementation or simply strategy performance evaluation. Most of the methodologies used in the literature on battery arbitrage bidding are variations of Model Predictive Control (MPC), Dynamic Programming (DP), Mixed Integer Linear Programming (MILP), Mixed Integer Non-Linear Programming (MINLP), and Reinforcement Learning (RL) models. The literature review in this paper aims to discuss these design- and methodological variations in detail.

For the single-market ID price arbitrage setting, DP methods are particularly suitable, considering that it is a dynamic system with time-dependency of states, sequential decision-making and identifiable state transition functions. A DP solution relies on Bellman's principle of optimality by breaking problems into smaller, overlapping sub-problems and solving them recursively, using backward induction, assigning value-to-go functions to each state on the basis that all remaining decisions have been optimised. This formulation excels in the presence of non-linearities and complex dependencies across states but suffers from the curse of dimensionality as computational cost increases exponentially with the number of states and horizon length, typically requiring discretisation and pruning methods for solvability. In their bi-level setup, (Baker et al., 2023) rely on a DP-based approach to solve a deterministic ID price-response problem based on historical data in the first stage to then use the resulting 'optimal opportunity value function' to train a Long-Short Term Memory (LSTM) model that tries to predict the optimal charge decision in each state given unseen future prices. This approach deviates from alternative stochastic approaches in that instead of maximising expected revenue given forecasted price scenarios, it chooses the best action based on the past 5 hours of actual price data and its previously trained relation to the pre-optimised value-to-go functions. The authors extend their work in (Alghumayjan et al., 2024), by introducing DA economic bidding in a two-settlement setup and using the previously described ID arbitrage strategy to maximise profits while considering the supply and demand deviations from the post-clearing DA schedules. The authors decide on the forecasted ID price as the optimal DA bid- and offer price in this setting.

However, the DP papers discussed so far do not directly model the price uncertainty of the ID bidding process. (Zheng et al., 2022) and (Abdulla et al., 2018) use Stochastic Dynamic Programming (SDP) methods to design strategies that can be used in an online setting and directly incorporate price uncertainty. (Zheng et al., 2022), also focusing on ID arbitraging, develop an analytical DP solution to derive the value function associated with price-SoC pairs and execute their control policy on a variety of trained Markov Price Processes. The different stochastic models are trained by varying the stage-dependency settings between price nodes, using both Real-Time (RT) prices directly and DA-RT differences as inputs, and experimenting with seasonal patterns. They achieve a 73% profit ratio compared to the perfect information case and find that DA-RT differences are more robust Markov Process inputs and that taking price autocorrelation into account improves performance by 10%. They also show that the granularity of SoC segmentation (created as part of the discretisation necessary for DP) has an optimal threshold, as a function of price fluctuation, such that profit performance of the DP approach is close to the fully flexible MILP solution in the Perfect Information case, with reportedly better computational efficiency. (Abdulla et al., 2018) solve a capacity constrained arbitrage problem – that also takes battery degradation into consideration – of a storage, co-located with a solar PV unit, by applying DP to point forecasts of the ID prices and the PV generation, and incorporate the results in a Model Predictive Control (MPC) framework. MPC is an inherently stochastic approach that – instead of solving for all actions at all stages at once – uses a model to predict future outcomes and solves an optimisation problem for a finite receding horizon at each time step, where only the first control action is implemented, after which the horizon rolls forward and the predictions are updated, and the system re-optimised. For RT or ID markets where most of the trades tend to happen in the final hour before delivery, so decisions are made in sequential, uncertain environments, MPC has been a widely applied optimisation method (see also (Arnold & Andersson, 2011)). (Abdulla et al., 2018) only generate a single scenario – without consideration of densities – for each interval and instead validate their results by sampling from multiple customers' solar PV generation against a Naïve Set-Point Control solution. Their findings reveal that even basic price forecasting and battery degradation models, like assigning a constant per-kWh cost to discharging, achieve good life-time revenue performance relative to Naïve benchmarks, though improving upon the price forecasting methods have more significant impact in bridging the earnings gap to the Full Information solution.

Though some studies neglect its consideration, battery degradation is a crucial element in realistic energy storage optimisation given that frequent cycling reduces the lifetime of a battery before it needs to be replaced. (Vetter et al., 2005) identify cycle depth, rate of charge, frequency of over-charging and average SoC as the main influences on cycle aging but highlight cycle depth as the far most critical in the case of a grid-scale unit. The nonlinear aging property is illustrated by the findings of (Ecker et al., 2014), in that a 7 Wh Lithium Nickel Manganese Cobalt Oxide (NMC) battery cell can perform around 50,000 cycles at 10% cycle depth over its lifetime, in contrast to only 500 cycles at 100% cycle depth, resulting in a reduction of a factor of 10 in total energy throughput. Therefore, accounting for cycle aging in a control optimisation requires appropriate cycle depth estimation, which is a non-trivial problem given SoC profiles with many local extrema (i.e. where the direction of the current changes). In their deterministic ID arbitrage optimisation, (Xu et al., 2018) produce a piecewise linearisation of the polynomial cycle depth stress function – which couldn't be directly incorporated into the MILP solution they use – by segmenting the SoC for the discharge portion of the cycles and using the results from a rain-flow counting algorithm as their ex-post benchmark, which counts cycles by iteratively combining them from smallest to largest. They assign a different cycling cost to a given discharge depending on which segment it fell into. Their results confirm that entirely disregarding cycle aging leads to negative profits, that constant per-KWh costs are robust base models since top spreads are worth a deep cycle, while the granularity of the SoC segmentation matters more when price deviations are smaller, i.e. in the constant-cost 1-segment-approximation case there are no activations for smaller arbitrage opportunities even though cycle aging cost would be low.

The papers discussed so far tended to focus on ID-only frameworks, however DA bidding is a fundamental component of most arbitrage strategies given its hedging ability. MILP is extensively used for optimised DA bidding in the literature, given the single-time market clearing for the full 24 hours, in contrast with the sequential nature of the ID market. In contract with DP formulations, MILP doesn't need additional fine tuning for discretisation as it can solve for continuous cases and produces accurate results efficiently with modern solvers. (Krishnamurthy et al., 2017) set up a DA bidding model for time domain arbitrage with the possibility of making corrections on the ID market, but without coordinating with the ID corrections during the DA stage, similar to (Alghumayjan et al., 2024). Instead of DP, however, they use Stochastic MILP formulations using Scenario Tree methods, Kernel Density Estimation (KDE) and Monte Carlo simulation models, after linearising the initially non-linear optimisation problem caused by the conditional bid acceptance indicator functions in some constraints. However, their solution seriously suffers from only imposing SoC constraints on the final power delivery values, but not the DA bid-offer values in their averaged price-quantity bidding optimisation, resulting in countless scenarios with infeasible post-clearing DA schedules that need to be corrected on the ID market at a significant, unaccounted cost. (Mohsenian-Rad, 2016) conduct an analysis that coordinates the bidding on the two markets, however they restrict the purpose of ID participation to only correcting the unmet or unsold energy from the DA-stage. Thereby, they avoid the issues of (Krishnamurthy et al., 2017), but eliminate the potential for additional ID arbitrage opportunities. They do, however, provide proof of their proposition that in their price-quantity bidding, storage operators need only a single segment in an optimal economic bid. In their model, the optimal price cutoff of the economic bid is argued to be a function of the dependency structure between DA and ID prices, i.e. given independence, ID price is found as an optimal cutoff, but given dependence they also condition it on expected DA-ID spread. They also provide sensitivity analysis on varying the cycle limit, battery efficiency and charge and discharge rates of the storage asset.

Finally, in recent years, due to its natural fit – given the goal of time-dependent cumulative reward maximisation in a dynamic system – RL has been more extensively explored in the literature for energy arbitrage, especially for ID markets. Highlighting the benefit of their distribution-free approach, (Wang & Zhang, 2018) present a basic Q-learning algorithm with epsilon-greedy exploration, proposing a mean-standardised reward function, although they completely disregard degradation and only consider maximum allowable charging rates. (Cao et al., 2020) improve upon this by feeding LSTM-CNN forecasted prices into a Deep Q-network that incorporates a battery degradation model and show that the RL solution improves on the SP baseline by 58% in their case study. Notably, however, they restrict the ID market to an RT framework without the possibility of transacting on multiple future settlement periods concurrently.

3. Methodology

This paper develops and evaluates bidding strategies given price uncertainties, with a focus on the multi-market arbitrage opportunities of storage operation by exploiting the price differences across time and across markets. For this analysis, focus is kept on the DA and ID markets. For the comprehensive evaluation of the online strategies, fully optimised offline solutions are derived as benchmarks. In an offline solution, future clearing price data – that could only be forecasted in a real-world application – is available at the time of optimisation, yielding the theoretical maximum value achievable for the specified objective function. By contrast, online optimisations involve decision-making under future uncertainty. Furthermore, based on the learnings from the literature, a constant degradation cost[1] – which was found to yield similar lifetime performances as more advanced cycle models – is built into the optimisations such that charge prices are increased and discharge prices are decreased by a constant amount. Accordingly, the cycling amount of the storage is also a linear sum of the charges, without differentiation between cycle depths.

[1] Also the modelling choice in (Alghumayjan et al., 2024) and (He et al., 2016).

Offline solutions are developed for each market separately as well as for several multi-market formulations, which allow non-physical trading across markets. In non-physical trading, an operator can buy and sell the same physical delivery contract (e.g. 10MWh energy delivery between 10AM and 11AM) at different times on different markets, earning on the price difference, and only has to deliver the net amount. Hence, prior volume commitments can be offset in subsequent bidding rounds. For the purposes of standardisation, daily cycle limits are imposed in the optimisations.

The multi-market approaches considered are as follows:

1) The **Simultaneous approach** assumes full knowledge of both DA and ID prices at the time of the DA bidding and solves for the optimal charge and discharge levels of both markets simultaneously as though there was only a single stage of bidding with simultaneous access to both markets. Notably, it only requires the cycle-limit to hold for the net position and not after each actual stage. Therefore, this formulation is conceptually akin to the one used in (Krishnamurthy et al., 2017).
2) The **Sequential approach** first optimises the DA market as though it was a single market case, then subsequently solves for the optimal charges and discharges on the ID market taking into account the inherited position resulting from the first stage DA bidding. Importantly, the cycle limit constraints need to be respected after each stage.
3) The **Integrated approach** combines elements of the Sequential and Simultaneous approaches. It solves the DA and ID bidding problems in two stages, first optimising the DA market in isolation, enforcing the cycle limit, but then "claws back" charges or discharges based on future ID prices before determining the final DA position. Consequently, an integrated approach can leave the final DA positions intentionally open, anticipating better opportunities on the ID market to close them. However, it is better hedged than the simultaneous approach due to pre-selecting the most optimal DA charging and discharging actions and avoiding speculative positions beyond those. Thereby, the integrated approach is an upgrade of the model developed in (Mohsenian-Rad, 2016), which also took ID prices into account during DA bidding but restricted the ID market to only filling in for uncleared DA bids.

Taking bidding optimisations online has different modelling implications for the auction-based DA market and the order-book based ID market in the GB system. Simulating the online pay-as-clear DA auction requires submitting bid and offer curves – consisting of price-quantity segments – for each auction period separately. The cleared volume of the storage operator is determined by intersecting the market clearing price line with the bid-offer curves, as illustrated in **Figure 1**. Any cleared volume's physical delivery is a

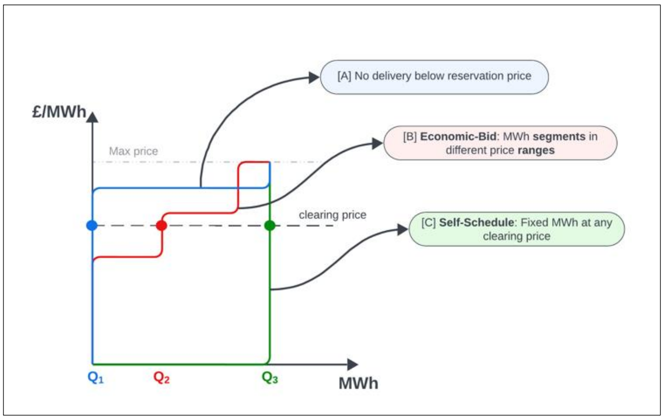

Figure 1. Offer clearing process on the Day-Ahead market, with examples of Economic Bidding and Self-Schedule bidding.

contractual obligation. The type of bid-offer curves that result in a fixed clearing volume regardless of the final clearing price are called Self-Schedule (or quantity-only) bids, in contrast with Economic (or price-quantity) bids that would lead to different volumes depending on the clearing price. This paper explores both Self-schedule and Economic bids for online DA bidding, though instead of making the parameters of the bid-offer curves subject to the optimisation in each settlement period – as done in (Krishnamurthy et al., 2017) – the bidding process is separated into first choosing which settlement periods to submit bids or offers for followed by determining the price segments.

Meanwhile, on the ID market, the order-book system allows for continuous trading up to 5 minutes before delivery, provided there is a counterparty willing to take the opposite position. Operating on a clear-as-bid basis, it results in multiple prices for the same delivery period. Further, the majority of the traded volume is in the final hour before physical delivery, as illustrated in (Modo Energy, 2023a) resembling an Hour-ahead Real-time bidding system, such as PJM in the U.S. Therefore, a fully realistic optimiser would contract ID positions on a rolling basis as the day progresses, while continuously updating its price forecasts, and executing multiple trades for the same delivery period depending on the size of its energy needs. Given this paper's scope, the online ID bidding process is simplified to a significant extent as the optimisation takes a position for all settlement periods at once at the start of the delivery day. Consequently, only Self-schedule bids are considered, since there is no possibility of subsequent correction of open positions. Despite such a significant simplification, the model is argued to serves as an appropriate lower bound for the achievable performance on the ID market for a smaller price-taking storage operator, since it relies on less information and has to make decisions earlier than in the real case, though it assumes that all intended trades will clear.

The online bidding strategies evaluated include 3 fully online and 2 hybrid approaches, as follows:

1) **Online DA bidding**: Self-schedule bids for the settlement periods when peaks and bottoms tended to be on average, in the previous 10 days.
2) **Online ID bidding**: Self-schedule bids based on ID price forecasts made with data available by the start of delivery day (such as the results of the DA market clearing).
3) **<u>Sequential</u>: PERFECT DA + Online ID:** As the fully offline sequential approach, first optimises the DA market as though it was a single market case, then optimises the ID charges and discharges, but using forecasted instead of actual ID prices.
4) **<u>Integrated</u>: Online DA + PERFECT ID:** Online Economic bidding on the DA market in the first stage based on ID forecasts available by DA gate closure, followed by position adjustments through the ID market based on perfect ID price information.
5) **<u>Integrated</u>: Online DA + Online ID:** Same as before, but ID bidding is based on forecasted ID prices.

This paper relies on Mixed Integer Linear Programming (MILP) formulations to optimise the arbitrage problems. MILP is an appropriate method for this analysis as it can fully accommodate the problem formulation and constraint system of the proposed approaches, relying on auxiliary variables and reformulations when necessary; it is straightforward to implement and interpret, can be solved in minimal time using professional solvers and doesn't involve compromises on precision or comprehensiveness as DP would, for example. Since forecast uncertainty in online cases is handled by back testing over many historical data points, rather than using simulated data from scenario trees, the benefits of the subproblem-based framework of DP[2] would also be redundant. Given the simplified model of the ID bidding process in this paper, computationally intensive methods such as RL or MPC, which excel in the continuous position updating formulation, are not needed. In all the MILP optimisations, the resolution of the optimisation is set at 30 minutes, the length of a Settlement Period in the GB markets, as of writing. The formulations further allow for specifically setting the duration, power capacity, cycle aging cost and cycle limit of the storage asset.

3.1. Offline Optimisation

3.1.1. Offline Optimisation: Single Market

a) DA-only and ID-only

In order to determine the optimal charge and discharge levels for a single market for each settlement period given full price information, the MILP aims to maximise the following objective function:

$$\sum_{t=1}^{T} q_{m,t,d} * d_{m,t} - q_{m,t,c} * c_{m,t}; \;\; m \in \{DA, ID\}, \tag{1}$$

subject to constraints (2)-(9):

$$0 \le c_{m,t} \le E_{max}, \qquad t = 1, \dots T; \;\; m \in \{DA, ID\}, \tag{2}$$

$$0 \le d_{m,t} \le E_{max}, \qquad t = 1, \dots T; \;\; m \in \{DA, ID\}, \tag{3}$$

$$SoC_t = SoC_{t-1} + c_{m,t} - d_{m,t}, \qquad t = 1, \dots T; \;\; m \in \{DA, ID\}, \tag{4}$$

$$0 \le SoC_t \le SoC_{limit}, \qquad t = 1, \dots T, \tag{5}$$

$$SoC_0 = SoC_T = 0, \tag{6}$$

$$cycle_0 = 0, \tag{7}$$

$$cycle_t = cycle_{t-1} + c_t, \qquad t = 1, \dots T, \tag{8}$$

$$cycle_T \le cycle_{limit}, \tag{9}$$

where parameters are defined by the following equations (10)-(16):

$$SoC_{limit} = P_{BESS} * D_{BESS} \tag{10}$$

$$E_{max} = \frac{P_{BESS}}{2} \tag{11}$$

$$cycle_{limit} = \#_{cycles} * SoC_{limit}, \tag{12}$$

[2] As seen in (Alghumayjan et al., 2024).

$$q_{DA,t,c} = p_{DA,t} + x_{BESS} \tag{13}$$

$$q_{DA,t,d} = p_{DA,t} - x_{BESS} \tag{14}$$

$$q_{ID,t,c} = p_{ID,t} + x_{BESS} \tag{15}$$

$$q_{ID,t,d} = p_{ID,t} - x_{BESS} \tag{16}$$

The objective function (1) sums up the values of the discharged energy across all the periods in the optimisation and subtracts the total costs of the charges. In the value determination, the prices are adjusted for degradation such that the charge price is offset in the positive, and the discharge price in the negative direction by a constant cycling cost as can be seen in (13)-(16). Constraints (2) and (3) ensure that the energy transfer in a given optimisation period doesn't exceed the storage's physical energy delivery capacities for neither charges nor discharges. Constraints (4)-(5) provide a way to track the SoC throughout the day and to ensure it is kept within operational limits, while (6) establishes that the initial and final States of Charge (SoC) of the storage must be 0, which eliminates the possibility of across day price arbitrages. The derivation of the upper SoC limit is provided in (10), multiplying the power rating of the storage by its duration. Finally, constraints (7)-(9) keep track of the total charge delivered since the start of day – which in this paper is the proxy for the amount of cycling – and they ensure that it does not exceed the cycling limit, which is a function of the upper SoC limit, and the number of cycles allowed, as expressed in (12).

3.1.2. Offline Optimisation: Multi-market

In the multi-market case where both DA and ID charge and discharge levels are subject to optimisation, the objective function is adjusted such that the MILP aims to maximise the sum of the discharge revenues from both markets minus the sum of the charging costs from both markets across the day:

$$\sum_{t=1}^{T} q_{DA,t,d} * d_{t,DA} + q_{ID,t,d} * d_{t,ID} - q_{DA,t,c} * c_{t,DA} - q_{ID,t,c} * c_{t,ID} \tag{17}$$

Further, constraints (2)-(4) are substituted by constraints (18)-(21):

$$0 \le c_{m,t}, \qquad t = 1, \dots T;\ \ m \in \{DA, ID\}, \tag{18}$$

$$0 \le d_{m,t}, \qquad t = 1, \dots T;\ \ m \in \{DA, ID\}, \tag{19}$$

$$-E_{max} \le c_{DA,t} + \ c_{ID,t} - d_{DA,t} - \ d_{ID,t} \le E_{max}, \qquad t = 1, \dots T, \tag{20}$$

$$SoC_t = SoC_{t-1} + c_{DA,t} + c_{ID,t} - d_{DA,t} - c_{ID,t}, \qquad t = 1, \dots T, \tag{21}$$

The substitution reflects that when both the DA and ID markets are active, it is the netted values of all the contracted charge and discharge levels that will determine the final physical delivery of the storage for a given delivery period. Constraints (2)-(3), which previously restricted the charge and discharge levels – on only a single market – to be non-negative values below the maximum energy transfer limit, are now replaced by (18)-(20). (18)-(19) uphold the non-negativity requirement but remove the upper limit on a single market's energy transfer due to non-physical trading, where multiple contracts in opposite directions can exceed individual battery limits but remain within operational constraints when combined. Instead, (20) now enforces upper limits, ensuring that net discharges or charges never exceed physical capacity. Meanwhile, constraint (21) replaces (4), indicating that the SoC of the battery at period t will change by the netted sum of the charges and discharges from the two markets.

b) DA + ID Simultaneous

The first multi-market optimisation this paper develops is the Simultaneous approach. In addition to the previously introduced constraints (18)-(21), the simultaneous trading formulation further requires setting the following constraint (22):

$$c_{DA,t} + c_{ID,t} + d_{DA,t} + d_{ID,t} \leq E_{max} * 2, \qquad t = 1, \dots T, \tag{22}$$

Constraint (22) is necessary, because if only the final netted total charge and discharge levels were constrained as done in (20) then there would be a possibility of carrying out profit-guaranteed non-physical trades with infinitely big volume sizes of opposite directions leading to infinite profits, which the MILP couldn't handle. The limit on the absolute sum of charges and discharges from both markets is set to two times the maximum energy transfer capacity, in a discretionary manner, as this allows for taking a full (i.e. maximum energy transfer) position in one direction on one market and for that position to be fully cancelled on the other market, leaving a net zero position overall.

Furthermore, given the non-physical trading allowance in the present and subsequent multi-market formulations, additional optimisation constraints need to be introduced so that the cycle limit can be appropriately upheld. The cycling in this paper's methodology is traced by summing over the physical charges of the battery throughout the day, which – in case of non-physical trading – is complicated by charges on one market being offset on another, leading to lower actual physical charge levels. To properly account for actual charges, it is the netted value of charges and discharges that needs to be traced, conditional on the netted value being positive. However, such an "if" condition introduces non-linearities that cannot be handled by a MILP solver. This problem can be circumvented by a reformulation that this paper introduces, using parameter M, which is a sufficiently large constant, auxiliary binary variable, δ_t, and the following constraints (23)-(28):

$$net_t = c_{DA,t} + c_{ID,t} - d_{DA,t} - d_{ID,t} \tag{23}$$

$$ac_t \geq net_t, \qquad t = 1, \dots T \tag{24}$$

$$ac_t \geq 0, \qquad t = 1, \dots T \tag{25}$$

$$ac_t \leq net_t + M * (1 - \delta_t), \qquad t = 1, \dots T \tag{26}$$

$$ac_t \leq M * \delta_t, \qquad t = 1, \dots T \tag{27}$$

$$cycle_t = cycle_{t-1} + ac_t, \qquad t = 1, \dots T, \tag{28}$$

$$M = 1000, \qquad \delta \in \{0,1\}$$

Constraint (23) introduces the net charge, net_t, as the difference between the sum of charges and sum of discharges in a given period across both markets. Constraints (24)-(27) introduce ac_t, the actual charge variable, and ensure, as proposed by this paper, that the optimised value of ac_t is 0 if the net_t is negative and is equal to net_t if it is non-negative. The proof follows from observing that if net_t is positive then (24) and (27) can only be satisfied concurrently if $\delta_t = 1$ holds, as it makes (27) a loose constraint given the size of large M. If $\delta_t = 1$, then (26) simplifies to $ac_t \leq net_t$ and together with (24) necessitates $ac_t = net_t$. Alternatively, if net_t is negative then (25) and (26) can only both be satisfied if $\delta_t = 0$, otherwise ac_t would need to be smaller than a negative number and bigger than 0 at the same time, which is infeasible. Therefore, if $\delta_t = 0$, then (27) simplifies to $ac_t \leq 0$, and together with (25) means that $0 \leq ac_t \leq 0$, which is equivalent to $ac_t = 0$. It has thus been shown that when the net charge is positive, the actual charge is equal to the net charge, and when the net charge is negative, the actual charge is set to 0 (when $net_t = 0$, then $\delta_t = 1 \; or \; 0$ lead to the same results). It follows that (28) appropriately tracks the cycling rate of the battery throughout the day, by summing up the actual charges that the storage is required to carry out once all the trades have been settled, and thereby replaces constraint (8), previously used for single market cases.

c) DA + ID <u>Sequential</u>

The sequential bidding approach with non-physical trading differs from the simultaneous approach in that it first optimises the DA market as though it was a single market with a cycling limit, and then reoptimizes the position on the ID market while the same cycle limit needs to be satisfied by the final netted physical positions. Therefore, in the first stage the single market MILP setup with constraints (2)-(14) can be run on the DA price data, while in the second stage the same constraint system can be applied as in the multi-market simultaneous approach but with the additional constraints (29)-(30) that bind the DA charge and discharge values to the outputs of the first stage optimisation:

$$c_{DA,t} = c_{stage1,t}, \qquad t = 1, \dots T, \tag{29}$$

$$d_{DA,t} = d_{stage1,t}, \qquad t = 1, \dots T, \tag{30}$$

d) DA + ID <u>Integrated</u>

The integrated approach operates essentially as a clawback mechanism of the sequential approach. After the first stage optimisation of the DA market, it observes the ID prices and retrospectively drops DA charges and discharges if it deems it optimal. In a MILP formulation, this can be achieved by first optimising the DA market in isolation, then using the same constraint system as the sequential approach in the second stage but, unlike (29)-(30), loosening the additional constraints (31)-(32) which require the DA actions to be no more than their first stage outputs. This allows the solver to reset previous DA charges and discharges to zero if that is more optimal, while it disallows non-zero DA trades that weren't carried out in the first stage.

$$c_{DA,t} \leq c_{stage1,t}, \qquad t = 1, \dots T, \tag{31}$$

$$d_{DA,t} \leq d_{stage1,t}, \qquad t = 1, \dots T, \tag{32}$$

3.2. Online Optimisation:

3.2.1. Online DA optimisation

The first online optimisation model that this paper develops is for single market DA bidding. As previously explained, the DA bidding process involves submitting segmented bid-offer curves, which reveal the actual contracted volume once the uniform clearing price has been determined for the entire market. For typical generators, like a gas plants, segmented bid curves are essential due to their varying marginal operating costs, which usually have an increasing return to scale property. Therefore, at a higher clearing price they find it profitable to increase operations to a higher marginal cost level. However, in the case of a storage system, for which the cycle aging cost could be considered the operating cost, multiple studies[3] have demonstrated that a constant cycle aging cost, – which would mean constant returns to scale – is a robust approximation of the power function describing the true stress function. Therefore, this paper chooses to focus on bid-offer curves with a single segment where the quantity splits are $0\ MWh$ and E_{max}, while acknowledging that results would be more precise with an accurate degradation cost accounting, likely leading to partial charges.

Since any bid that was cleared in the DA auction is a binding contract, it is not beneficial to repeat bids across multiple periods hoping to secure just one in the presence of uncertainty, as multiple clearances could lead to commitments beyond the intended volume. Two potential approaches are bid splitting to smooth the volatility, or to try to identify the times with the highest likelihood of having a local price extremum and bid only then. This paper, given its scope, focuses on the latter. Therefore, an identification methodology is required to find the peak and bottom times. Fortunately, the outputs of the offline optimisation of the DA market contain the peak and bottom timing information required, which we will use to inform the online DA bidding.

a) Self-schedule, Rolling Average timing DA bidding

A Self-schedule bidding strategy is proposed for the DA-only approach, which eliminates the possibility of unclosed positions due to uncleared bids. The timing of the charges and discharges in this strategy are determined as a rolling average of the previous 7 days' peak and bottom times, identified from the offline stage. Hyperparameter tuning was conducted to determine the best rolling window.

3.2.2. Online ID optimisation

b) Self-Schedule ID-only bidding

The online ID-only models developed in this paper are simple Stochastic Programming solutions without recourse, using the same constraints (2)-(9) as the offline case but applied on forecasted prices.

3.2.3. Online multi-market optimisation

This paper elaborates on three distinct online multi-market optimisation frameworks, while varying which of the two markets is treated with price uncertainty.

[3] (Xu et al., 2018)

c) Sequential, PERFECT DA + Online ID

The first model replicates the sequential multi-market model developed for offline prices, but treating the ID prices as uncertain, making the second stage a single-scenario SP formulation.

d) Integrated, Online DA + PERFECT ID

The second online multi-market model is a variation of the offline integrated model, however the previously applied clawback method upon observing future ID prices is no longer possible due to taking the DA bidding online. In a real-world online DA scenario, there is no mechanism to reverse DA positions nor is there perfect information on future ID prices. Therefore, the DA positioning needs to be optimised using the tools of bid-offer curve segmentation at the storage optimiser's disposal. Upon deeper consideration, it becomes clear that the DA position clawback in the offline case happens when the ID price is lower than the DA price while the storage asset was scheduled to charge, or when the ID price is higher than the DA price while the storage asset was scheduled to discharge. DA actions clawed back this way can be replaced with the same volume on the ID market at a better price, while DA actions spared this way can simply be offset on the ID market at a pure profit, if the optimiser chooses to do so. In an auction framework, this would effectively translate to setting the – previously elaborated – single segment price cutoff of both the bid and the offer curves to the ID price. Hence, a bid volume wouldn't clear if the same energy could be bought at a lower price on the ID market, and the same logic applies with offers. Since the actual ID price isn't available during online DA bidding, this paper uses the forecasted ID price as its cutoff. In the second stage, despite the online DA bidding, this model uses an offline ID correction based on actual ID prices.

e) Integrated, Online DA + Online ID

Finally, the third multi-market model rests on the same foundation as the second one, but the second stage ID correction is now also online, based on forecasted ID prices.

3.2.4. Intra-Day Price Forecasting

Multiple ID price forecasts are compared in this paper. First is a Naïve benchmark forecast that takes the average of the previous 10 days' prices for that given time of day. A separation of weekends and weekdays was also tried, but results showed that the daily autoregressive nature of the price data is more significant than the difference between the weekend and weekday price shapes. The main forecasting method this paper relies on uses a Gradient Boosted Decision Tree (GBT) model, where the features are the price points at the same time of day from the previous 10 days', in addition to indicator variables for Saturday, Sunday and Monday (the days of and after the weekend). A second GBT model is also trained with the inclusion of the cleared DA prices, which can be used during the ID stage of the online bidding strategies, when this information would be available in the real world as well. The GBT models were trained using the XGBoost package, optimising the Root Mean Squared Errors (RMSE). Following along the seminal book from (Hastie, 2001), basic regularisation steps were taken by setting the maximum tree size to 3, the learning rate to 0.1 and establishing an early stopping rule of 10. The DA price was the most important feature by measure of number of splits – when it was available – followed by the price from the previous day. This is still a relatively basic model that doesn't include variables such as weather and temperature forecasts, interconnector usage or demand expectations, however, given that the focus of this paper is on the benchmarking of bidding strategies, the GBT models prove to be sufficient bases. More advanced forecasting models can be developed as an extension and be easily plugged into the system built throughout this paper.

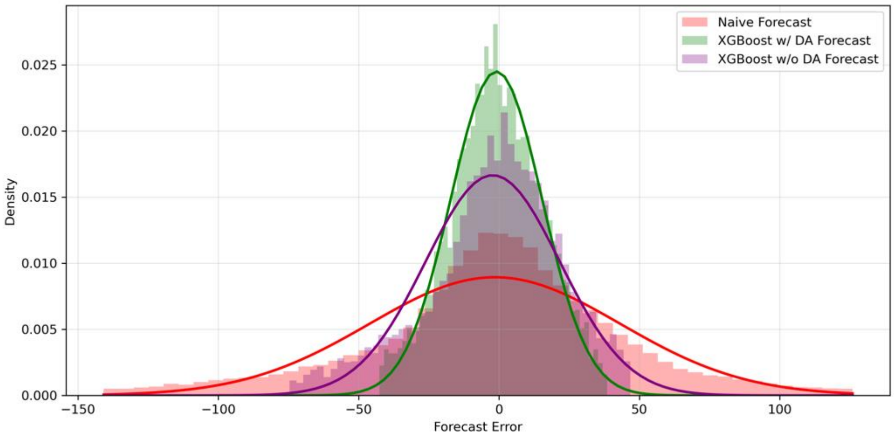

Figure 2. Standardised middle 95% percentile distributions of forecast errors of Naïve and XGBoost forecasts (with and without DA prices as features) of ID prices, with Normal Distribution approximations.

The performance of the different ID forecasting models is reported in **Figure 2**, as standardised density distributions of forecast errors, further approximated by normal distributions. The Naïve model has a standard error of 44£/MWh in the test period, while adding the cleared DA prices as a feature improves the standard errors of the XGBoost model from 25£/MWh to 16£/MWh.

4. Data

The hourly and half-hourly DA price data was taken from Great Britain's N2EX wholesale electricity market for the available period of January 1st, 2020, to December 31st, 2024. The ID price data used in the modelling is the reference price index for half-hour product trades in the order book (RPD HH) for each half-hour period of the day from the EPEX electricity market, aggregated at the end of the trading day, collected from July 1st, 2021, to December 31st, 2024. Both series were collected using Modo Energy's API access point. There was a total of 36 missing values out of 58,465 in the ID dataset, spread across the 2,5-year period with no missing period longer than 2 hours. Therefore, a simple linear interpolation was used to fill in missing values for sake of full continuity but impact on overall forecasting and bidding results is arguably negligible.

Figure 3 presents the prices for a sample period. It is visually apparent that the half-hourly DA auction prices exhibit a co-integration relationship with the hourly clearing prices with a strong error correction mechanism, while the ID prices can deviate significantly and tend to be more volatile, though there are days when all three series are closely aligned. Given its smoothness, which is more conducive to accurately identifying the timing of the peaks and bottoms, the hourly DA price was chosen as the reference price for the DA market, disaggregated to half-hourly resolution using linear interpolation. The XGBoost ID forecasting model was trained and validated on a combined dataset of 2021 and 2024 prices, as they each resembled market dynamics in 2022 and 2023, respectively – as shown in **Figure 4** –, preserving a 2-year long coherent testing and simulation period in the middle. This 2-year period between January 1st, 2022, and December 31st, 2023, is an ideal natural experiment setting to evaluate bidding strategies – and their adaptability – as it includes periods of extreme as well as normal price environments. Energy prices spiked in early 2022 after international sanctions were imposed on Russian fossil fuels due to its invasion of Ukraine in February 2022, and Russia shut down the Nord Stream gas pipeline causing European countries competing to find alternative energy sources in a supply constrained environment. Meanwhile, by 2023, given newly diversified portfolios, a milder winter and market corrections after an apparent overreaction, prices normalised to pre-crisis levels.

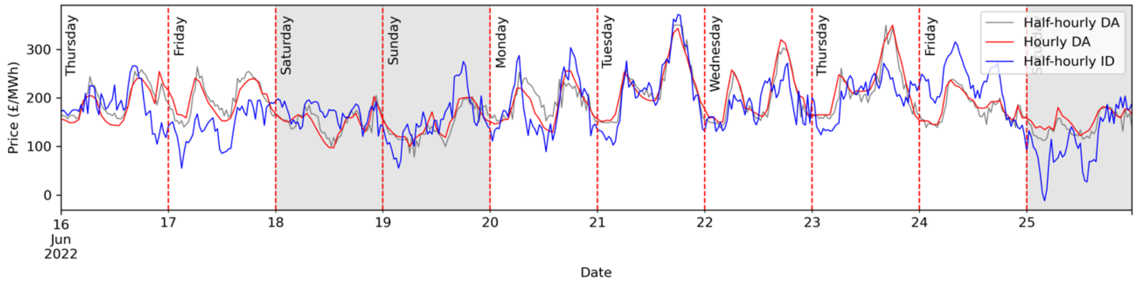

Figure 3. Half-hourly price curves of Hourly and Half-hourly DA - and Half-hourly ID markets for a sample period in June 2022, with days separated and weekends identified by shading.

Additionally, we can note from **Figure 4 (b)** that daily spreads are strongly correlated with price levels while it also confirms that they are consistently higher on the ID market than the DA market, especially in 2022.

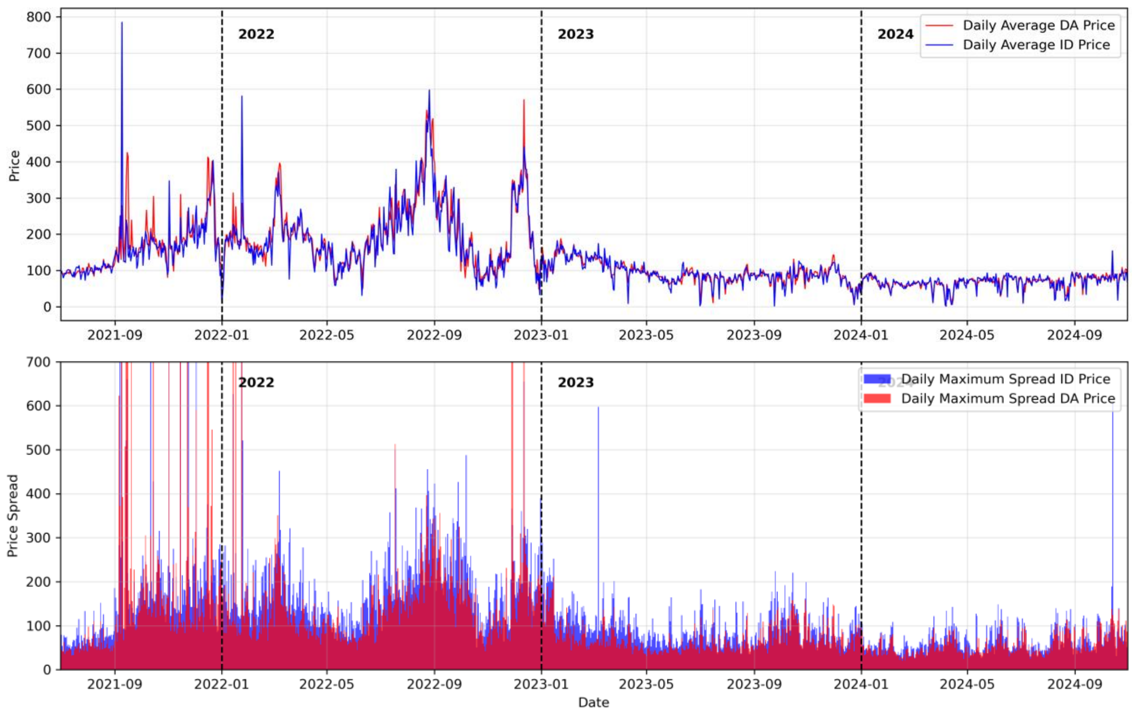

Figure 4. Daily average prices **(a)** and daily maximum price spreads **(b)** on DA and ID markets, from 2021 June to 2024 December.

5. Results

5.1. Offline Optimisation Results

A large 100MW/100MWh storage system with a cycle limit of 2 and degradation cost of 5£/MWh – which is half of what is used in (Alghumayjan et al., 2024) who only assign costs to discharging – is used as the basis of the simulations. **Figure 5 (a)-(e)** present the results of the offline optimisation for a sample day of February 20th, 2022, serving as a case study to ascertain the validity of the MILP method and to observe the dynamics of arbitrage trading. The results reveal that, for all the bidding strategies, the cycle limit is respected, SoC constraints are maintained, and the net charge never exceeds the storage's half hourly energy transfer capacity. From a strategic perspective, charges and discharges happen at bottom and peak times, respectively, as expected, while the integrated approach in **(e)** appropriately drops DA charges and discharges from the first stage when the ID price is lower or higher, respectively, at the given time. It takes the GLPK solver less than 0.1 seconds to run a single strategy's optimisation for a day, and about 60 seconds to run 2 years of simulated results. The code for the integrated strategy is presented in the Appendix.

The DA-only optimisation results in **Figure 5 (a)** show that the premium of running a 2nd cycle on the DA market is the price difference between the first discharge and the second charge, which is significantly less than the returns from running a single cycle of charging around 5AM and discharging at 6PM. The ID-only results in **Figure 5 (b)** highlight that the price shape on the ID market can deviate substantially from the typical "2-hump shape" on the DA market, leading to different optimised charging times. While the higher price variability means more uncertainty about the timing of the local extrema, it also results in a 2-cycle-premium that is comparable to a single cycle's value, and a total profit that is 66% higher than with the DA-only strategy.

Figure 5 (c) presents the results of the simultaneous optimisation of DA and ID markets. Non-physical trades driven by intermarket arbitrage opportunities occur in almost every period based on which market's price was momentarily higher. Some physical charge actions, mainly on the ID market, still exploit the intertemporal price differences when they exceed the intermarket differences within-period, though only 1.5 cycles are conducted in total. With the exception of trading the most extreme intertemporal arbitrage opportunities, the simultaneous approach does not consider overall price shapes, focusing heavily on the momentary relative spread between DA and ID prices, and thereby consistently commits to DA charging even during periods of daily price peaks. Similar to (Krishnamurthy et al., 2017), it also results in DA schedules that are severely infeasible on their own from a SoC perspective. While this works in an offline setting, it is risky in an online setting where the DA-ID spread is difficult to forecast, and there is no guarantee that the DA commitments can be profitably offset on the ID market while ensuring that the storage's SoC constraints are met.

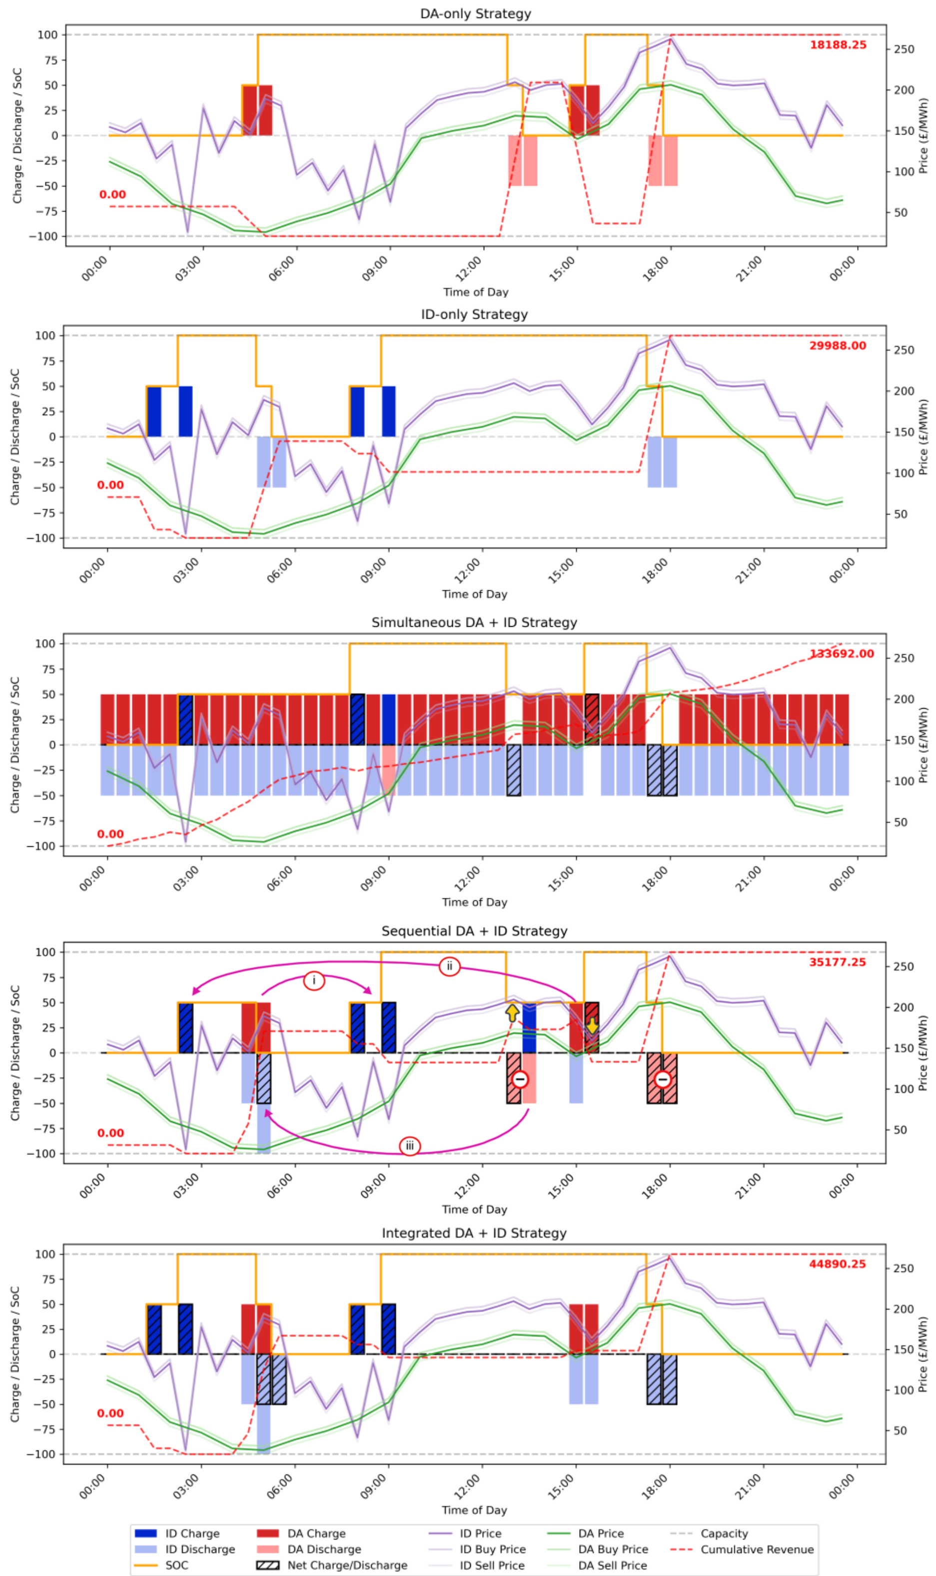

Figure 5. Optimised battery operation of different strategies for the sample day February 20th, 2022. **(a)** DA-only strategy **(b)** ID-only strategy **(c)** Simultaneous strategy **(d)** Sequential strategy **(e)** Integrated strategy. Charge, Discharge and Net Charge levels from the two markets along with State of Charge levels are visualised for each half-hour, with value index to the left. DA and ID prices visualised along right axis. Cumulative Revenue is tracked throughout along a hidden third axis.

The sequential and integrated strategies in **Figure 5 (d) and (e)** highlight the fundamental logic of ID position updating. Since there is usually a full 2-cycle DA commitment in the first stage, ID corrections must come in pairs to ensure that the operational day ends at zero SoC – given the decision to not include the Balancing Mechanism in the model –, i.e. any DA position closed or opened by an ID action has to be matched with another ID action in the opposite direction that opens or closes another position. The monetary impact of the correction is the difference between ID prices prevalent at the timing of the two actions. The types of ID corrections can be broadly categorised into three groups: i) there is a pure cross-market arbitrage opportunity in closing a prior DA position, ii) there is an ID position in a different period at a better price than the existing DA position, which needs to be offset, iii) an ID opportunity is not feasible due to SoC limitations imposed by a DA position, which needs to be "shifted across time". The sample day studied in **(d)** presents examples of all three. Type i) appears when the optimiser offsets 5AM DA charges while ID prices were much lower – generating significant non-physical profits – and replaces it with ID charges at 8AM and 9AM at slightly higher charging prices. Meanwhile, moving the charge from 3PM to 2:30AM offers better prices, as described by type ii). However, such an offset would make the previously created 9AM charge later on impossible due to exceeding SoC constraints. Therefore, offsetting the DA discharge at 1:30PM and moving it to 5:30AM on the ID market accommodates the earlier trade-off at negligible offset costs, demonstrating the type iii) correction.

The advantages of the integrated approach in **(e)** are highlighted by comparing it to the sequential case where both DA discharge actions happened while ID prices were higher. In the sequential approach, the 1:30PM DA discharge was offset at a negative cost, while the 1PM discharge couldn't be offset due to the unfavourable ID price differentials at the time, indicated by the arrows. The integrated approach, avoiding such offset costs by "clawing back" prior DA discharge commitments, manages to carry out two ID charges before 3AM, compared to just one. In addition, the 5:30PM DA discharges previously blocked the possibility of exploiting the higher ID prices, as the integrated approach does. Meanwhile, DA charges were retained in the integration, as they could be offset at a pure profit due to higher ID prices at the time. In summary, the integrated approach pre-emptively eliminates DA discharges that would block or deteriorate ID opportunities while it also exploits inter-market price arbitrage opportunities on the remaining DA positions. Overall, the difference in approach results in a revenue increase of 29% on the sample day.

The relative daily and cumulative performance of the offline strategies is visually represented in **Figure 6 (a)-(c)**. **(a)** and **(b)** show the daily profits for February and March in 2022 and 2023, respectively, with the previously discussed sample day, February 20th, 2022, highlighted in **(a)**. The figures reveal that the ID-only approach consistently outperforms DA-only, robust to changing market conditions, while the performance of the sequential approach relative to the ID-only approach changes substantially from 2022 to 2023. In 2022, the superiority of the two strategies' alternates from day-to-day, with the sequential strategy more steadily exceeding the ID-only strategy towards the end of the year as seen in **(c)**. However, this reverts in 2023 as price volatility settles down and there are fewer non-physical trades to capitalise on from early DA commitments. The logic and development of the integrated strategy is validated by the observable strict dominance of its performance over the single-market and sequential strategies. As expected, the simultaneous strategy outperforms the rest of the group by a factor of 3 on average, though the degree of differentiation is irregular from day-to-day, indicating that ID prices are not just simple error correction processes around the DA price. Overall, the cumulative returns curve flattened out for all strategies in 2023 as prices normalised after a more volatile year in 2022 and returns proved to be stable with occasional spikes. Over the entire period of 2 years, for the studied 1-hour, 2-cycle storage, ID-only and sequential strategies outperformed the DA-only strategy by 26%, while the integrated approach outperformed them both by an additional 33%.

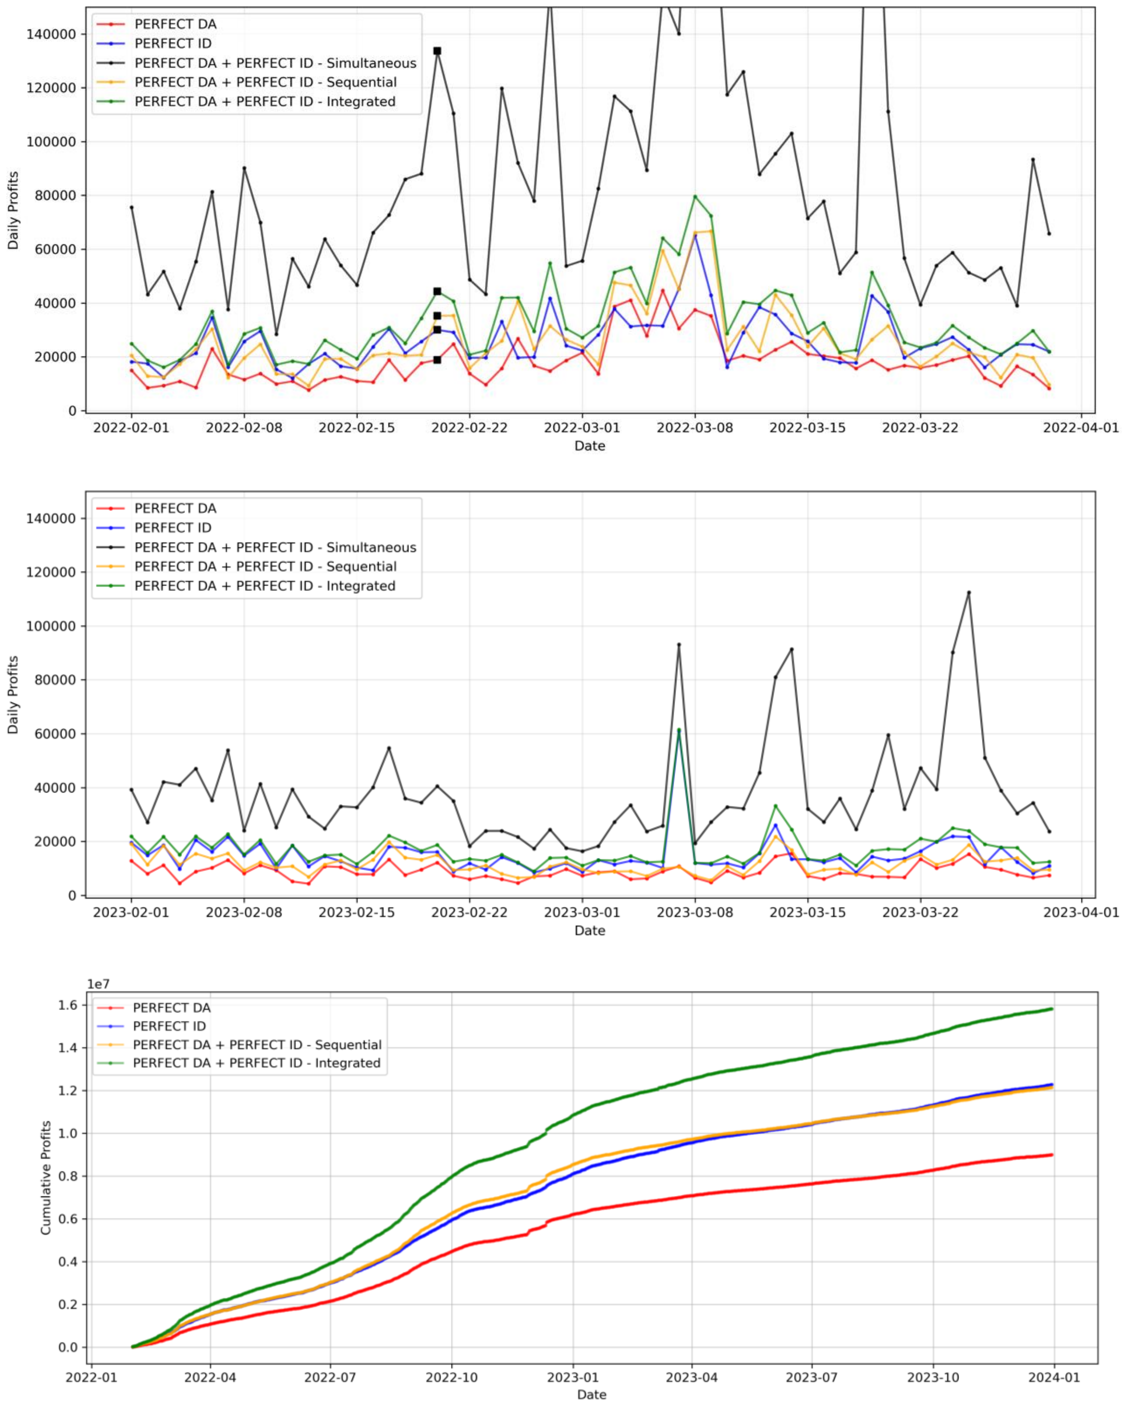

Figure 6. Daily profits for offline DA-only, ID-only, Simultaneous, Sequential and Integrated strategies for February and March in 2022 **(a)** and 2023 **(b)**, with case study sample day of February 20th, 2022, highlighted. Cumulative profits of strategies for the entire period of 2022-2023 in **(c)**.

5.2. Online Optimisation Results

Figure 8 presents the times of day of the offline optimised charges and discharges for the year 2023, which confirm the high degree of autoregressive nature as well as the seasonality of the data, which underlies the online DA Self-schedule and Economic bidding methods that this paper develops. It also shows that the charge and discharge times are significantly more sporadic on weekends. **Figure 7** contrasts the performance of the online Self-schedule optimisation with that of the offline version, disaggregated by day of the week, further proving that it is more difficult to forecast the timing of the peaks and bottoms on weekends, which would require additional fine tuning. Nonetheless, the online DA optimisation achieves 81% of the total potential profits in 2023.

Figure 8. Timing of offline optimised Charges and Discharges on the DA market for each day of 2023, with differentiation between Weekdays and Weekends.

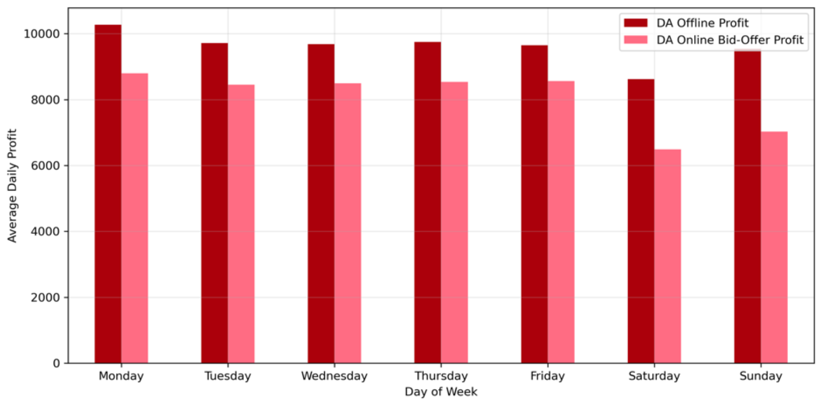

Figure 7. Relative performance of Offline and Online DA bidding approaches, disaggregated by day of the week, for the year 2023.

Figure 9 (a) and **(b)** summarise the relative performance of the series of strategies developed throughout the paper, relying on Naïve and GBT forecasting methods in case of online components, respectively. The plot distinguishes between DA-only, ID-only, Sequential and Integrated strategies with the inclusion of price uncertainty varied across the two markets, allowing us to determine the impact that each modelling modification has on the overall performance and capture rate, i.e. the ratio of profits that an online approach can achieve relative to the offline perfect information solution.

As partially reported earlier, the Self-schedule DA-only strategy, which is the same across the two plots as it doesn't rely on ID forecasts, achieves a capture rate of 73% in 2022 and 81% in 2023, using the rolling average peak-timing approach. The ID-only strategy captures merely 55% of profits for the entire two-year period when relying on Naïve forecasts, pushing it below the online DA-only performance, due to its higher volatility, though it achieves a superior 82% capture rate using GBT price forecasts, which exceeds the capture rate of 73% found in (Zheng et al., 2022). In general, it is observed that using weak forecasts during online ID-adjustment such as in the semi-online sequential approach or the fully online integrated approach, leads to worse returns than not participating on the ID market at all. However, the fully online integrated strategy using stronger GBT forecasts manages to beat both online single-market approaches, achieving a capture rate of 73%.

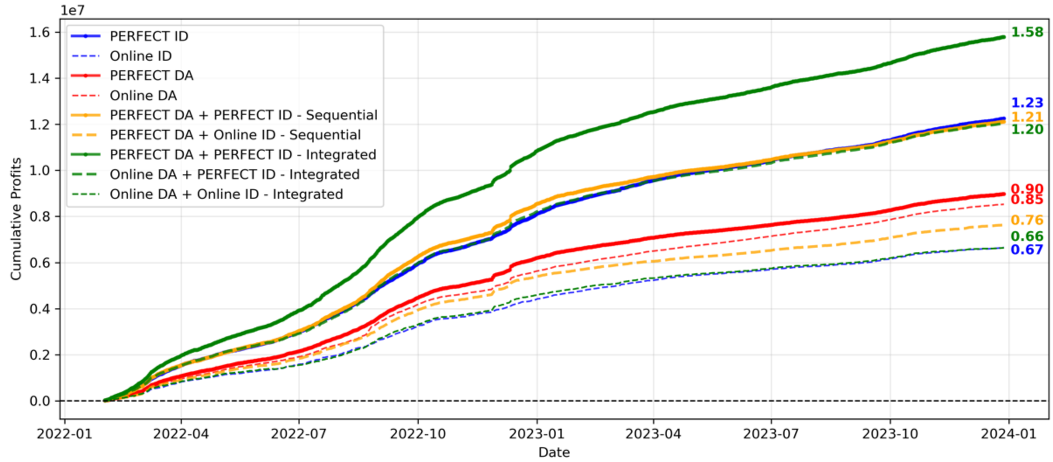

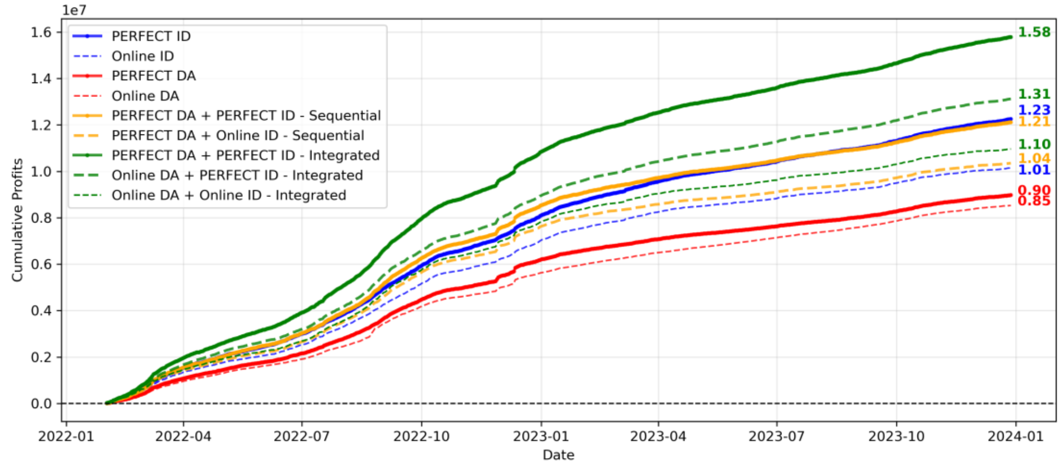

Figure 9. Cumulative profits of Offline and Online Strategies. Underlying Naïve **(a)** and Gradient Boosted Tree **(b)** methods used for ID price forecasting in the Online formulations are in separate figures. Colours differentiate between strategy types, dashed lines indicate the presence of modelling uncertainty, while thickness shows the presence of some degree of Perfect Information.

Further, while the online integrated approaches based on the Naïve forecasts trend very closely with the corresponding ID-only approaches, and fail to surpass the sequential approaches, using the GBT forecasts accentuate the hedging ability of the integrated approach, as intended: it takes more positions on the DA market when the sequential approach performs well, however, when spreads on the ID market are more favourable and the sequential approach slightly flattens out – in 2023 in particular – the integrated approach takes fewer DA commitments and tracks the performance of ID-only more closely, thereby displaying a form of weak dominance over both alternative approaches. This is most apparent from the semi-online integrated strategy's performance surpassing even fully offline alternative approaches. These results suggest that the GBT forecast generated before DA market clearing manages to estimate the DA-ID spread direction correctly more often than random, indicating that markets are not fully efficient.

5.3. Sensitivity Analysis

The paper so far has operated with a 100MW/100MWh, 2-cycle limit assumption, however as an investor there is a wider range of storage specifications available with potentially different applications and returns on investment. Therefore, a series of scenarios have been simulated for the developed offline strategies of DA-only, ID-only and the Integrated approach, by varying the storage's duration between 1, 2, and 4 hours and its cycle limit from 1, 2, to 3 cycles, while keeping the degradation cost fixed. The simulated results for the year 2022 are presented in **Figure 10 (a) and (b)**. Increasing the duration of the battery from 1 to 2, and from 2 to 4 hours increases potential revenues of the integrated approach by 78%, and 58% respectively, more than for the single-market approaches. It is not clear though how much of this can be captured using online methods, given that the top spread opportunities are seized first and it is expected that higher bidding accuracy is required to exploit tighter spreads. Nonetheless, 98% of newly installed batteries in Q4 2024 were 2 hours or longer in duration[4], supporting the results. **Figure 10 (b)** confirms

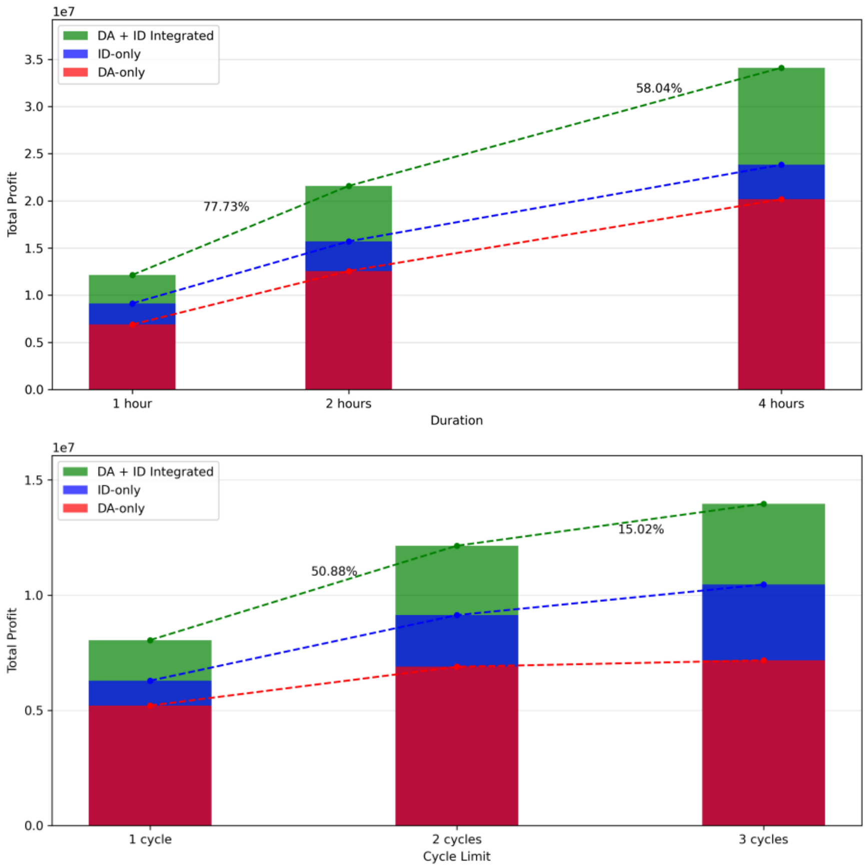

Figure 10. Performance – for 2022 – of ID-only, DA-only and Integrated strategies given variation in battery duration between 1,2 and 4 hours while fixing daily cycle limit at 2 cycles **(a)** and variation in cycle limit from 1, 2 to 3 cycles while fixing duration at 1 hour **(b)**.

[4] Modo Energy, 2024

what we've seen for a single day earlier, that an additional cycle has higher premium on the ID market than on the DA market, and a 3rd cycle – taking degradation costs into consideration – yields negligible uplift on the DA market. An additional cycle is significant in an integrated strategy at a 51% premium but returns are quite limited beyond that. Notably, the same sensitivity analysis yields almost identical results when simulated for 2023, despite the significantly different market conditions previously observed, which highlights that although profits are affected, the underlying anatomy of the arbitrage trading is fairly stable. However, as the penetration of PV solar generation continues and the "duck curve" becomes more pronounced, the premiums from additional cycles are expected to increase as well.

6. Conclusion

This paper developed an array of arbitrage trading strategies for the Day-Ahead (DA) and Intra-Day (ID) markets, using MILP methods to find the optimal bidding prices and quantities. Offline solutions, based on perfect price information, were derived for the single market DA and ID cases, and for their multi-market combination in different configurations, i.e. a simultaneous approach under which both market's quantities are decided together as though their bidding was concurrent, a sequential approach of DA commitments in the first stage followed by ID corrections in the second stage such that DA bidding doesn't consider future ID conditions, and another sequential approach that integrates future ID information into its DA positioning but, unlike the simultaneous approach, still ensures that the DA position would be operationally and financially sensible on its own. A constant degradation cost as well as daily cycle limits are factored into the optimisation frameworks. The paper proposes a MILP reformulation for the multi-market frameworks that is able to handle the cycle limit constraints imposed on the physical net positions from combining the bids and offers of the two markets. An in-depth case study analysis of a sample day demonstrates the theory behind ID position correction and the underlying value in the integrated strategy. The second portion of the paper discusses the development of online versions of the strategies. An online Economic-bidding DA-only strategy composed of single-segment bid-offer curves with a cutoff price set at the forecasted ID price is explored. The order-book based ID bidding was simplified to a Self-schedule single-scenario Stochastic Programming formulation without recourse. Naïve and Gradient Boosted Tree forecasting models of ID prices were compared to evaluate the effect of different prediction errors. The performance of the strategies was evaluated on historical data. Results reporting daily and cumulative profits reveal that the ID-only and the multi-market Sequential strategy perform very similarly on average, both exceeding DA-only profits by about 26% while the Integrated strategy outperforms them both by an additional 33%. Though the performance of strategies using online bidding on the ID market are poorer than the online DA-only strategy's performance when using Naïve forecasts, a Gradient Boosted Tree model leads to significantly higher capture rates. The weak dominance of the semi-online Integrated strategy over the fully offline ID-only and Sequential strategies underpins its adaptability to changing market conditions. The analysis closes out with a sensitivity analysis of varying the duration and cycle limit used in the offline optimisations, finding significant uplifts from moving from 1- to 2- and to 4 hours in duration. The premium of a second cycle is also found to be substantial while a third cycle has much lower benefits.

6.1. Limitations and Future Work

There are a number of natural extensions to the work in this paper, as well as more substantial pieces of research that could build on its results. The most immediate limitation is the sophistication of the price forecasting methodology used which currently doesn't consider weather or demand variables nor supply factors such as carbon and gas prices or projected interconnector activity. In addition to feature extensions, the modelling method could be extended to methods that are able to incorporate more complex or structured interactions than the XGBoost model such as Deep Neural Networks or Markov Chains. Forecasts that consider the dependence structure of prices across markets could yield further benefits. Another extension would be to conduct a sensitivity analysis on not only offline formulations, but online versions as well since capture rates could deviate significantly as the duration of a battery increases.

More significant augmentations would include a non-linear degradation cost function, which better approximates the power functions that mathematically describe cycle aging. It would require assigning varying amounts of cycling costs depending on the depth of a given discharge and is expected to yield partial charges in an optimal solution. In an online setting, non-linear degradation costs would introduce added complexity as charge volumes would be a continuous function of spread forecasts. Furthermore, the current online DA bidding used the forecasted ID price as its cutoff, which is argued to be optimal when DA and ID prices are independent. However, if prices are dependent, then if the realised ID price is higher relative to its forecasted mean, the probability of the same for DA increases as well. Therefore, a bid-offer cutoff strategy that conditions the bid-size on the size of the DA-ID spreads directly could yield more robust results. More dynamic conditional bidding would also open up the possibility of deeper evaluation of the speculative approach as well, which lacked proper risk-management in the static strategy formulations explored in this paper. Finally, while this paper focused on the bidding strategies themselves, implementing them as they are would not be optimal from an investment perspective, due to lacking the larger context. Generating long-term price expectations and specifying risk preferences and investment horizons could guide the determination of the minimum spread sizes that a storage should even activate for.

More considerable future extensions could unbind the ID bidding simplifications and incorporate continuous position updating within an order book framework. Also, additional markets' integration – such as the Balancing Mechanism in a capacity splitting or the frequency response services in a capacity stacking framework – could yield outcomes closer to real-world storage operations.

7. References

Abdulla, K., De Hoog, J., Muenzel, V., Suits, F., Steer, K., Wirth, A., & Halgamuge, S. (2018). Optimal Operation of Energy Storage Systems Considering Forecasts and Battery Degradation. *IEEE Transactions on Smart Grid*, *9*(3), 2086–2096. https://doi.org/10.1109/TSG.2016.2606490

Alghumayjan, S., Han, J., Zheng, N., Yi, M., & Xu, B. (2024). Energy storage arbitrage in two-settlement markets: A transformer-based approach. *Electric Power Systems Research*, *235*. https://doi.org/10.1016/j.epsr.2024.110755

Arnold, M., & Andersson, G. (2011). *Model Predictive Control of Energy Storage including Uncertain Forecasts*.

Baker, Y., Zheng, N., & Xu, B. (2023). Transferable Energy Storage Bidder. *IEEE Transactions on Power Systems*, *39*, 4117–4126. https://api.semanticscholar.org/CorpusID:255393977

Cao, J., Harrold, D., Fan, Z., Morstyn, T., Healey, D., & Li, K. (2020). Deep Reinforcement Learning-Based Energy Storage Arbitrage with Accurate Lithium-Ion Battery Degradation Model. *IEEE Transactions on Smart Grid*, *11*(5), 4513–4521. https://doi.org/10.1109/TSG.2020.2986333

Department for Energy Security and Net Zero. (2024). *Energy Trends: December 2024*.

Ecker, M., Nieto, N., Käbitz, S., Schmalstieg, J., Blanke, H., Warnecke, A., & Sauer, D. U. (2014). Calendar and cycle life study of Li(NiMnCo)O2-based 18650 lithium-ion batteries. *Journal of Power Sources*, *248*, 839–851. https://api.semanticscholar.org/CorpusID:98721247

Hastie, T. and T. R. and F. J. (2001). *The elements of statistical learning: data mining, inference, and prediction*. Springer.

He, G., Chen, Q., Kang, C., Pinson, P., & Xia, Q. (2016). Optimal Bidding Strategy of Battery Storage in Power Markets Considering Performance-Based Regulation and Battery Cycle Life. *IEEE Transactions on Smart Grid*, *7*(5), 2359–2367. https://doi.org/10.1109/TSG.2015.2424314

Krishnamurthy, D., Uckun, C., Zhou, Z., Thimmapuram, P. R., & Botterud, A. (2017). Energy Storage Arbitrage Under Day-Ahead and Real-Time Price Uncertainty. *IEEE Transactions on Power Systems*, *33*(1), 84–93. https://doi.org/10.1109/tpwrs.2017.2685347

Modo Energy. (2023a). *Trading in the continuous intraday market: how does it work?*

Modo Energy. (2023b). *Wholesale trading: what are the different markets and how do they work?*

Modo Energy. (2024a). *GB BESS Outlook Q2 2024: Battery revenue stacking and dispatch optimization*.

Modo Energy. (2024b). *GB BESS Outlook Q2 2024: How will battery markets evolve?*

Mohsenian-Rad, H. (2016). Optimal bidding, scheduling, and deployment of battery systems in California day-ahead energy market. *IEEE Transactions on Power Systems*, *31*(1), 442–453. https://doi.org/10.1109/TPWRS.2015.2394355

Vetter, J., Novák, P., Wagner, M. R., Veit, C., Möller, K. C., Besenhard, J. O., Winter, M., Wohlfahrt-Mehrens, M., Vogler, C., & Hammouche, A. (2005). Ageing mechanisms in lithium-ion batteries. *Journal of Power Sources*, *147*(1–2), 269–281. https://doi.org/10.1016/j.jpowsour.2005.01.006

Wang, H., & Zhang, B. (2018). Energy Storage Arbitrage in Real-Time Markets via Reinforcement Learning. *2018 IEEE Power & Energy Society General Meeting (PESGM)*, 1–5. https://doi.org/10.1109/PESGM.2018.8586321

Xu, B., Zhao, J., Zheng, T., Litvinov, E., & Kirschen, D. S. (2018). Factoring the Cycle Aging Cost of Batteries Participating in Electricity Markets. *IEEE Transactions on Power Systems*, *33*(2), 2248–2259. https://doi.org/10.1109/TPWRS.2017.2733339

Zheng, N., Jaworski, J., & Xu, B. (2022). Arbitraging Variable Efficiency Energy Storage Using Analytical Stochastic Dynamic Programming. *IEEE Transactions on Power Systems*, *37*(6), 4785–4795. https://doi.org/10.1109/TPWRS.2022.3154353

8. Appendix

```
#INTEGRATED NON-PHYSICAL
#DA_cycle_limit = 2

def seq_nonph_clawback_daily_profits(date, combined_da_id, capacity, cycling_cost, cycle_limit, DA_cycle_limit):

#---------------------------------
# STEP 1: DA CHARGES PRE-CALCULATION
#---------------------------------
    optimised_times = pd.date_range(start=date, end=date + pd.Timedelta(days=1), freq='30Min', inclusive='left', tz='UTC')

    da_prices_optimisation = {}
    for time in optimised_times:
        da_prices_optimisation[time] = combined_da_id.loc[time]['DA_price']

    # MODEL SOLVING
    model = ConcreteModel()

    # SETS
    model.I = Set(initialize=optimised_times)

    # PARAMETERS
    model.max_capacity = Param(initialize=capacity / 2)
    model.DA_price = Param(model.I, initialize=da_prices_optimisation)
    model.charge_price = Param(model.I, initialize={i: price + cycling_cost for i, price in da_prices_optimisation.items()})
    model.discharge_price = Param(model.I, initialize={i: price - cycling_cost for i, price in da_prices_optimisation.items()})
    model.SOC_limit = Param(initialize=capacity * duration)
    model.cycle_limit = Param(initialize= DA_cycle_limit * capacity * duration)


    # VARIABLES
    model.charge = Var(model.I, within=NonNegativeReals, initialize=0)
    model.discharge = Var(model.I, within=NonNegativeReals, initialize=0)
    model.SOC = Var(model.I, within=NonNegativeReals, initialize=0)
    model.cycle = Var(model.I, within=NonNegativeReals, initialize=0)


    # CONSTRAINTS
    def charging_cap_upper(model, i):
        return model.charge[i] - model.discharge[i] <= model.max_capacity
    model.upper_capacity_constraint = Constraint(model.I, rule=charging_cap_upper)

    def charging_cap_lower(model, i):
        return model.charge[i] - model.discharge[i] >= -model.max_capacity
    model.lower_capacity_constraint = Constraint(model.I, rule=charging_cap_lower)

    def SOC_cumulation(model, i):
        if model.I.ord(i) == 1:
            return model.SOC[i] == model.charge[i] - model.discharge[i]
        else:
            return model.SOC[i] == model.SOC[model.I.prev(i)] + model.charge[i] - model.discharge[i]
```

```
    model.SOC_cumulation_constraint = Constraint(model.I, rule=SOC_cumulation)

    def SOC_constraint(model, i):
        return model.SOC[i] <= model.SOC_limit
    model.SOC_limit_constraint = Constraint(model.I, rule=SOC_constraint)

    def cycle_cumulation(model, i):
        if model.I.ord(i) == 1:
            return model.cycle[i] == model.charge[i]
        else:
            return model.cycle[i] == model.cycle[model.I.prev(i)] + model.charge[i]
    model.cycle_cumulation_constraint = Constraint(model.I, rule=cycle_cumulation)

    def cycling_constraint(model, i):
        return model.cycle[i] <= model.cycle_limit
    model.cycle_limit_constraint = Constraint(model.I, rule=cycling_constraint)

    # OBJECTIVE FUNCTION
    def objective_function(model):
        return sum((model.discharge_price[i] * model.discharge[i] - model.charge_price[i] * model.charge[i]) for i in model.I)
    model.objective = Objective(rule=objective_function, sense=maximize)

    # SOLVING THE MODEL
    solver = SolverFactory('glpk')
    results = solver.solve(model)

    # SAVING THE RESULTS
    DA_charges = {}
    for i in model.I:
        DA_charges[i] = model.charge[i].value

    DA_discharges = {}
    for i in model.I:
        DA_discharges[i] = model.discharge[i].value

#-----------------------------------
# STEP 2: ID CORRECTION
#-----------------------------------
    optimised_times = pd.date_range(start=date, end=date + pd.Timedelta(days=1), freq='30Min', inclusive='left', tz='UTC')

    id_prices_optimisation = {}
    for time in optimised_times:
        price = combined_da_id.loc[time]['RPD_HH_price']
        if pd.isna(price):
            # Return None if any price is NaN
            return None, None, None, None, None, None, None, None, None, None
        id_prices_optimisation[time] = price

    da_prices_optimisation = {}
    for time in optimised_times:
        da_prices_optimisation[time] = combined_da_id.loc[time]['DA_price']
```

```
model = ConcreteModel()

# SETS
model.I = Set(initialize = optimised_times)

# PAREMETERS
model.max_capacity = Param(initialize = capacity / 2) # Divide by 2 due to optimisation happening on a half-hourly scale
model.ID_price = Param(model.I, initialize = id_prices_optimisation)
model.ID_charge_price = Param(model.I, initialize = {i: price + cycling_cost for i, price in id_prices_optimisation.items()})
model.ID_discharge_price = Param(model.I, initialize = {i: price - cycling_cost for i, price in id_prices_optimisation.items()})
model.DA_price = Param(model.I, initialize = da_prices_optimisation)
model.DA_charge_price = Param(model.I, initialize = {i: price + cycling_cost for i, price in da_prices_optimisation.items()})
model.DA_discharge_price = Param(model.I, initialize = {i: price - cycling_cost for i, price in da_prices_optimisation.items()})
model.SOC_limit = Param(initialize = capacity * duration)
model.cycle_limit = Param(initialize= cycle_limit * capacity)
model.DA_charge_limit = Param(model.I, initialize = {i: charge for i, charge in DA_charges.items()})
model.DA_discharge_limit = Param(model.I, initialize = {i: discharge for i, discharge in DA_discharges.items()})


# VARIABLES
model.DA_charge = Var(model.I, within = NonNegativeReals, initialize=0)
model.DA_discharge = Var(model.I, within = NonNegativeReals, initialize = 0)
model.ID_charge = Var(model.I, within = NonNegativeReals, initialize=0)
model.ID_discharge = Var(model.I, within = NonNegativeReals, initialize = 0)
model.SOC = Var(model.I, within = NonNegativeReals, initialize = 0)
model.cycle = Var(model.I, within=NonNegativeReals, initialize=0)
model.net = Var(model.I, within=Reals, initialize = 0)

# AUXILIARY VARIABLES
model.actual_charge = Var(model.I, within=NonNegativeReals)
model.delta = Var(model.I, within=Binary)
M = 1000


def charging_cap_upper(model, i):
    return model.DA_charge[i] + model.ID_charge[i] <= model.max_capacity
model.upper_capacity_constraint = Constraint(model.I, rule = charging_cap_upper)

def charging_cap_lower(model, i):
    return model.DA_discharge[i] + model.ID_discharge[i] <= model.max_capacity
model.lower_capacity_constraint = Constraint(model.I, rule = charging_cap_lower)

def da_charge_clawback(model, i):
    return model.DA_charge[i] <= model.DA_charge_limit[i]
model.da_charge_clawback = Constraint(model.I, rule = da_charge_clawback)

def da_discharge_clawback(model, i):
    return model.DA_discharge[i] <= model.DA_discharge_limit[i]
model.da_discharge_clawback = Constraint(model.I, rule = da_discharge_clawback)
```

```
def SOC_cumulation(model, i):
    #print(model.I.ord(pd.to_datetime('2022-03-23 00:00:00+00:00')))
    if model.I.ord(i) == 1:
        return model.SOC[i] == model.DA_charge[i] - model.DA_discharge[i] + model.ID_charge[i] - model.ID_discharge[i]
    else:
        return model.SOC[i] == model.SOC[model.I.prev(i)] + model.DA_charge[i] - model.DA_discharge[i] + model.ID_charge[i] - model.ID_discharge[i]
model.SOC_cumulation_constraint= Constraint(model.I, rule = SOC_cumulation)

def SOC_constraint(model, i):
    return model.SOC[i] <= model.SOC_limit
model.SOC_limit_constraint = Constraint(model.I, rule = SOC_constraint)



#CONSTRAINTS ACCOUNTING FOR THE ID CANCELLING OF DA ACTIONS DURING CYCLE ACCUMULATION
def net_charge(model, i):
    return model.net[i] == model.DA_charge[i] + model.ID_charge[i] - model.DA_discharge[i] - model.ID_discharge[i]
model.net_charge_constraint = Constraint(model.I, rule = net_charge)

def ac_lower_bound_rule(model, i):
    return model.actual_charge[i] >= model.net[i]
model.AC_LB = Constraint(model.I, rule=ac_lower_bound_rule)

def ac_upper_bound_net_rule(model, i):
    return model.actual_charge[i] <= model.net[i] + M*(1 - model.delta[i])
model.AC_UB_net = Constraint(model.I, rule=ac_upper_bound_net_rule)

def ac_upper_bound_delta_rule(model, i):
    return model.actual_charge[i] <= M*model.delta[i]
model.AC_UB_delta = Constraint(model.I, rule=ac_upper_bound_delta_rule)

def cycle_cumulation(model, i):
    if model.I.ord(i) == 1:
        return model.cycle[i] == model.actual_charge[i]
    else:
        return model.cycle[i] == model.cycle[model.I.prev(i)] + model.actual_charge[i]
model.cycle_cumulation_constraint = Constraint(model.I, rule=cycle_cumulation)

def cycling_constraint(model, i):
    return model.cycle[i] <= model.cycle_limit
model.cycle_limit_constraint = Constraint(model.I, rule=cycling_constraint)



# OBJECTIVE FUNCTION
def objective_function(model):
    return sum((model.ID_discharge_price[i] * model.ID_discharge[i] + model.DA_discharge_price[i] * model.DA_discharge[i] - model.ID_charge_price[i] *
model.ID_charge[i] - model.DA_charge_price[i] * model.DA_charge[i]) for i in model.I)
model.objective = Objective(rule = objective_function, sense = maximize)

# SOLVING THE MODEL
solver = SolverFactory('glpk')
results = solver.solve(model)
```

```
    # Check if the solver found an optimal solution
    if results.solver.termination_condition != TerminationCondition.optimal:
        return None, None, None, None, None, None  # Return None if optimization is unsuccessful

    # Calculate the total profit for the day
    total_profit = sum((model.ID_discharge_price[i] * model.ID_discharge[i].value + model.DA_discharge_price[i] * model.DA_discharge[i].value -
model.ID_charge_price[i] * model.ID_charge[i].value - model.DA_charge_price[i] * model.DA_charge[i].value) for i in model.I)
    total_cycles = sum(model.ID_charge[i].value + model.DA_charge[i].value for i in model.I)
    ID_charge_prices = {i: model.ID_price[i] for i in model.I if model.ID_charge[i].value > 0}
    DA_charge_prices = {i: model.DA_price[i] for i in model.I if model.DA_charge[i].value > 0}
    ID_charge_times = {i: i for i in model.I if model.ID_charge[i].value > 0}
    DA_charge_times = {i: i for i in model.I if model.DA_charge[i].value > 0}
    ID_discharge_prices = {i: model.ID_price[i] for i in model.I if model.ID_discharge[i].value > 0}
    DA_discharge_prices = {i: model.DA_price[i] for i in model.I if model.DA_discharge[i].value > 0}
    ID_discharge_times = {i: i for i in model.I if model.ID_discharge[i].value > 0}
    DA_discharge_times = {i: i for i in model.I if model.DA_discharge[i].value > 0}

    return total_profit, total_cycles, ID_charge_prices, DA_charge_prices, ID_charge_times, DA_charge_times, ID_discharge_prices, DA_discharge_prices,
ID_discharge_times, DA_discharge_times

integrated_profits = []
for date in dates:
    profit, _, _, _, _, _, _, _, _, _ = seq_nonph_clawback_daily_profits(date, combined_da_id, capacity, cycling_cost, cycle_limit, DA_cycle_limit)
    if profit is None:
        integrated_profits.append(0)  # Use 0 for missing data
    else:
        integrated_profits.append(profit)
```