UNIVERSIDAD PONTIFICIA COMILLAS

ESCUELA TÉCNICA SUPERIOR DE INGENIERÍA (ICAI)

OFFICIAL MASTER'S DEGREE IN THE
ELECTRIC POWER INDUSTRY

Master's Thesis

Order Analysis and Design in OMIE's Intraday Markets

Author: Sebastián Maravall Almqvist
Supervisor: Nuria Trancho
Co-Supervisor: -

Madrid, July 2026

Master's Thesis Presentation Authorization

THE STUDENT:

NAME OF THE STUDENT

SEBASTIÁN MARAVALL ALMQVIST

THE SUPERVISOR

NURIA TRANCHO TEJERINA

Signed: Date: 17/07/2026

Authorization of the Master's Thesis Coordinator

Dr. Luis Olmos Camacho

Signed: ………………….. Date: ……/ ……/ ……

Abstract

The increasing penetration of variable renewable generation in the Spanish power system has increased the importance of the role of intraday auctions, yet OMIE's intraday order-type catalogue (simple orders and simple block orders) was assembled around the operating constraints of the thermal and pumped-hydro assets that have historically dominated it. As battery storage begins to scale toward Spain's PNIEC targets, it is not yet established whether that catalogue serves a cycling battery's operating profile or merely tolerates it. This thesis develops an automated backtesting framework that quantifies, in euros rather than in principle, what OMIE's existing order types deliver to a battery and what a proposed instrument modelled on what a Loop Block Order (C88) could add.

The methodology combines an empirical characterisation of how OMIE's own intraday auctions currently work and how the orders within them and concretely the block order is used in practice, parsed directly from a full year of published offer data, with a day-ahead and three-session intraday backtest of a 5 MW/10 MWh battery, in which an LP-optimised Simple order and a rule-based, forecast-driven Loop order are evaluated each session against a perfect-foresight welfare benchmark and against each other. The results show that OMIE's block order is used almost exclusively by combined-cycle gas turbines and pumped hydro, which together supply 87.4 % of all block submissions, while nuclear, wind, and solar on fewer than 0.1% of the thousands of intraday offers they submit, confirming the instrument as a targeted response to a specific physical constraint rather than a general-purpose product participants merely under-use. Against the perfect-foresight benchmark, a battery bidding the best available order type each session captures 81.4 % of achievable welfare across 387 operative days, with the residual gap attributable almost entirely to forecast error rather than order-type constraint.

Loop, evaluated as a policy counterfactual not currently offered by OMIE, wins outright on only 4.6–7.6 % of days across sessions and cannot match Simple's period-by-period flexibility, since most of the session-days require more legs than Loop's two-leg structure can express; its genuine value instead lies in partial downside mitigation, reducing or eliminating losses on 15 of Simple's 18 Session 1loss days. Attaching OMIE's own Minimum Income Condition logic to Loop as a reject-on-loss clearing rule recovers a further 2,740 € (0.7 %) on top of Loop's own realised revenue, increasing its "risk-mitigating" characteristic.

Keywords: electricity markets, day-ahead and intraday markets, OMIE, battery energy storage, complex orders, block orders, Loop Block Order, welfare-gap decomposition, order-type optimisation, linear programming

List of Acronyms

Table 0.1: List of acronyms used in the report

Acronym	Definition
ACER	Agency for the Cooperation of Energy Regulators
BESS	Battery Energy Storage System
C88	Loop Order (Loop Block Order)
CCGT	Combined-Cycle Gas Turbine
CNMC	Comisión Nacional de los Mercados y la Competencia (Spanish National Commission on Markets and Competition)
CVaR	Conditional Value at Risk
DA	Day-Ahead
EU	European Union
IDA	Intraday Auction
LASSO	Least Absolute Shrinkage and Selection Operator
LCOS	Levelized Cost of Storage
LEAR	Lasso Estimated AutoRegressive
LP	Linear Programme
MAE	Mean Absolute Error
MAV	Minimum Acceptance Volume
NECP	National Energy and Climate Plan
NEMO	Nominated Electricity Market Operator
OMIE	Operador del Mercado Ibérico de Energía
PHS	Pumped Hydro Storage
PNIEC	Plan Nacional Integrado de Energía y Clima
PV	Photovoltaic
RQ	Research Question
SCO	Scalable Complex Order
SDAC	Single Day-Ahead Coupling
SIDC	Single Intraday Coupling
SoC	State of Charge
TSO	Transmission System Operator
VaR	Value at Risk

Contents

1 | Introduction and Motivation

1.1 | Introduction to the European Electricity Market and NEMOs

The European electricity market is an interconnected system designed to harmonize national exchanges, couple the price among bidding zones, and align energy transactions with the physical reality of cross-border electricity flows [1] [2]. This integration is primarily enabled by the Capacity Allocation and Congestion Management (CACM) Regulation (EU) 2015/1222, which establishes the regulatory foundation for the Single Day-Ahead Coupling (SDAC) and the Single Intraday Coupling (SIDC) [3] [4]. These mechanisms allow electricity to be traded efficiently within countries and across national borders [1].

At the center of this framework are the **Nominated Electricity Market Operators (NEMOs)**: entities designated by national regulatory authorities to organize and operate electricity markets in coordination with other European operators [4]. They are the ones responsible for the organization of the European day-ahead market, which is the cornerstone of short-term electricity trading in Europe. It determines electricity prices for delivery on the following day through a single pan-European auction running once per day. Market participants submit their buy and sell orders, which are matched using the EUPHEMIA algorithm to simultaneously calculate clearing prices for all bidding zones and optimally allocate cross-border transmission capacity [5].

Figure 1.1: Map with European NEMOs

The European intraday market operates after the day-ahead session and is the main venue for adjusting supply and demand positions closer to real-time delivery. Its role has grown with rising renewable penetration: solar photovoltaic output responds to irradiance shifts within minutes, and wind generation carries forecast deviations that may only surface hours before delivery. Intraday rebalancing is not confined to renewable generators, however. A generation shortfall in any technology, whether caused by forecast error, an unexpected outage, or a technical constraint, requires offsetting supply from dispatchable resources such as combined-cycle gas turbines or reservoir hydro, both of which are regular intraday participants. The session draws from the full generation mix.

The intraday market consists of two complementary trading mechanisms: **continuous intraday market (XBID)** and **intraday auctions (IDAs)**, with the latter having been introduced across all European bidding zones on June 13th, 2024 [6]. The continuous market is a real-time shared order book operated through a different platform to the day-ahead. In it, participants can buy and sell electricity contracts until shortly before delivery (one hour or 30 minutes before delivery), which provides flexibility. Orders entered in one bidding zone can match with orders in any other, provided transmission capacity is available. Regular matching in the XBID system is triggered by the entry of a new order and follows price-time priority: orders are ranked by price (buy orders from highest to lowest, sell orders from lowest to highest) with submission time as the tiebreaker among orders at the same price level, and the trade price is always set by the passive order already present in the book.

However, intraday auctions work in the same way as the day-ahead market (they also go through the EUPHEMIA clearing mechanism), but they differ because they take place at different times of the day (closer and closer to the actual delivery) and that there are three (instead of only one) scheduled rounds. The first two come some hours after the day-ahead (1st at 15:00 D-1, 2nd at 22:00 D-1), while the third takes place in the morning of the same day the energy is sent (10:00 D). Prices are thus more explicit and auctions help increase liquidity for specific times throughout the day. Both mechanisms improve system flexibility and allow market participants to respond more effectively to forecast updates, unexpected outages, and renewable generation variability.

Despite the geographical diversity of European electricity markets, the operation of NEMOs is highly harmonized through common governance structures (such as the NEMO Committee[7]), methodologies, and technical systems defined by the CACM regulation [6]. This harmonization ensures efficient market functioning, non-discriminatory access, and consistent price formation across interconnected bidding zones. Rather than operating independently, NEMOs coordinate through multilateral cooperation frameworks to ensure the integrated operation of European electricity markets. The key elements of this harmonized framework include:

- **Common methodologies:** The All NEMO Committee facilitates cooperation among NEMOs in developing and implementing the common terms, conditions. Methodologies required under CACM ensure consistent rules for operation across Europe.
- **Unified pricing algorithm:** All NEMOs use the single price coupling algorithm EUPHEMIA for day-ahead and intraday auctions, ensuring consistent price calculation and optimal cross-border energy allocation across the coupled market.
- **Joint governance:** NEMOs and Transmission System Operators (TSOs) jointly oversee market coupling operations through governance structures such as the Market Coupling Steering Committee, enabling coordinated decision-making and operational efficiency. This unites the market clearing with the actual transmission of electricity, and assures it goes hand in hand without disparities in the energy dispatching.

- **Harmonized operational procedures:** NEMOs maintain shared procedures to ensure that market participant orders are matched according to common European rules, guaranteeing transparency, fairness, and market integrity.

Although both the European day-ahead and intraday markets operate under a harmonized regulatory framework defined by SIDC, the differences arise in how individual NEMOs and bidding zones implement specific trading products and operational rules. Here it is worth mentioning ACER, the Agency for the Cooperation of Energy Regulators, which is a key EU agency established to enhance cooperation among national energy regulatory authorities and ensure a well-functioning internal energy market in Europe. ACER is in charge of providing a product catalog and trading rulebook so European power markets can connect and trade seamlessly. It offers both mandatory and optional products, and the difference in how European NEMOs apply these shows divergences that arise due to enabling a broader range of order types, clearing structures, and bidding strategies throughout Europe [6] [1]. As a result, despite the common market coupling architecture, the practical design and availability of trading products vary significantly between regions. This thus structures our analysis around two main dimensions:

The first is the variation in intraday auction product design across NEMOs. Although the European intraday framework sets a common regulatory floor, order type portfolios vary across NEMOs in ways that shape how different technologies can participate. OMIE, GME, OTE, Nord Pool, and EPEX SPOT have each made distinct choices about which optional instruments to offer (block orders, complex conditions, linked bids, etc. which will be discussed later on) and these choices are not neutral. A market that lacks the instruments suited to a given technology's operational constraints effectively excludes it from efficient participation. Understanding where these portfolios diverge, and against what generation mix backdrop, is a precondition for evaluating whether current designs serve the resources now entering European grids.

The second is the impact these bidding and order structures have on market outcomes. Product availability is only half the question. The other half is whether complex order types are actually used, by whom, and to what effect on price formation, liquidity, and arbitrage between auction and continuous trading. OMIE is the only NEMO among those analysed that publishes offer-level data (quantity, price, and order type), subject to a 90-day confidentiality period, for both day-ahead and intraday sessions, making it the empirical anchor for this dimension. The cross-NEMO comparison remains, but grounded in market design and regulatory evidence rather than transaction data other operators do not disclose.

These two dimensions frame the analysis that follows: what each technology can bid, and whether it actually does.

1.2 | Motivation

The intraday market's role as a flexibility mechanism, described above, materialises in practice through the products NEMOs make available to participants. Before analysing how any market clears, a participant or researcher must first map the order types in play: what positions they allow, under what conditions they fill, and how the auction aggregates them into a price. The day-ahead market establishes this vocabulary through simple orders, block orders, and complex conditions. The intraday auctions inherit and extend it, but not uniformly. EPEX SPOT, GME, Nord Pool, OTE, and OMIE each offer a distinct product catalogue for both their day-ahead and intraday sessions. The CACM regulation harmonizes the architecture, but it does not harmonize the instruments.

For flexibility resources, this divergence is not abstract. A battery storage operator designing a bidding strategy for the Spanish electricity market must work within OMIE's specific rulebook, which differs from the products available in neighbouring bidding zones. As BESS deployment accelerates across Europe and Spain, storage assets that could arbitrage price spreads across auction windows can face constraints set not just by physics or economics, but also by the products available in the catalogue their local NEMO offers. Quantifying that constraint requires a model that links order type selection to realised revenues. This thesis builds a model for OMIE that evaluates how different order structures (present in other NEMOs) can affect BESS bidding performance in the Spanish market.

The starting point for that analysis is a precise account of the order types themselves: the instruments through which participants express their physical and economic constraints in each market. It is only after having performed a thorough analysis of the order that the evaluation on whether some of these could be used to fit batteries can be performed.

1.3 | Order types

Within the following analysis, the different order types will be explained. ACER's product catalogue defines the full set of mandatory and optional order types available under the CACM framework; each NEMO then selects which optional instruments to implement for its own participants. All NEMOs now have to offer MTU orders of 15 minutes, and most of them also offer block orders alongside the mandatory simple orders. The rest depend on each NEMO and what their market participants require to reflect their needs and restrictions as closely as possible to reality.

1.3.1 | Simple (MTU) Orders

The first order type is the simple order, formally termed an aggregated period order in EUPHEMIA documentation. This order type specifies a single price-quantity pair for one Market Time Unit (MTU). Each such order involves four elements: quantity, price, time period, and the bidding zone. Sell orders from all participants in a bidding zone are sorted from lowest to highest price and aggregated into a single supply curve; purchase orders are sorted from highest to lowest. Orders clearing in-the-money must be fully accepted; out-of-the-money orders must be fully rejected; at-the-money orders may be partially accepted or rejected. [5]

The MTU defines the temporal granularity of the order. EUPHEMIA supports periods of 15, 30, or 60 minutes, however, the official clearing price is now calculated on a 15-minute basis. Participants can still submit 60 and 30-minute orders, but EUPHEMIA, since the 1st October 2025, disaggregates them into four 15-minute sub-periods for clearing [8]. The MTU is 15', but the order granularity available to participants is still flexible [5].

Simple orders are the only mandatory order type across all NEMOs. All other instruments are optional and implemented at each NEMO's discretion.

Category	Product / Condition	Technical Logic and Description
MTU Orders	**15', 30' and 60'**	Market Time Units where 15' is the new base standard. Prices for 30' and 60' units are derived using the "Average Rule," which is the strict arithmetic average of the underlying 15' clearing prices.
	Aggregated Curves	Supply and demand offers aggregated into single curves (linear, stepwise, or hybrid) per bidding zone.

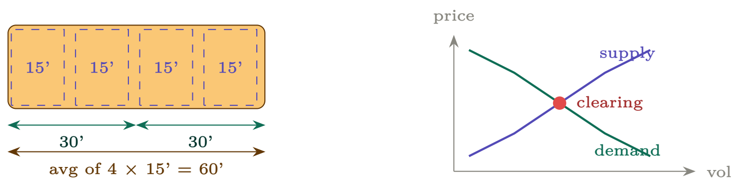

Left: 30' and 60' blocks are averages of their underlying 15-min EUPHEMIA slots. *Right*: Supply/demand curves; intersection = clearing price.

1.3.2 | Block Orders

Block orders encompass several hours at the same price and a fixed volume profile and are thus executed at the same ratio across all hours. They act as a single indivisible unit composed of a contiguous span of multiple MTUs. Block orders operate on a fill-or-kill basis; they are accepted or rejected in their entirety. Acceptance is conditional on the volume-weighted average of clearing prices across the block's delivery periods meeting or exceeding the submitted limit price (when it is used to buy).

They can also be used as a minimum acceptance ratio (a selling block), setting a floor on the share of the order that must be matched for execution to proceed — a feature motivated by fixed costs such as plant start-up expenses. Block orders are optional under EUPHEMIA; each Market Operator decides whether to offer them. However, all the NEMOs analysed in the thesis (and in general, most NEMOS) do. It is the most commonly used complex order and the base upon which the rest of these order types are built.

Category	Product / Condition	Technical Logic and Description
Block Orders	**Standard Block**	An order with a fixed price limit and volume for a specific set of MTUs. It is accepted based on the volume-weighted average price across those periods.
	Profile Block	A block order where the volume can vary for each MTU within the block, allowing for customized generation profiles.

Left: Standard block — fixed price and volume across all MTUs. *Right*: Profile block — volume differs per MTU.

1.3.3 | Complex Order Blocks

The first smart blocks to be described are linked block orders. These are families of blocks structured around a parent-child hierarchy, where the execution of each child order is conditional on the acceptance of its parent. In the standard two-block configuration, the parent can be accepted independently, but the child cannot be accepted without the parent first being accepted; however, the child can contribute surplus sufficient to make the parent economically viable, effectively saving the parent from rejection. The acceptance ratio of the parent must always be greater than or equal to those of all its child orders, and multiple levels of hierarchy are permitted, though the number of children and levels is constrained by each exchange. This structure allows generators to represent a range of production profiles across different price scenarios, submitting a base-load block as parent and incremental output blocks as children. [9]

Exclusive groups are collections of block orders within which at most one may be accepted. The sum of acceptance ratios across all orders in the group is constrained to be no greater than one, so that if the minimum acceptance ratio is set to one for each order, the structure reduces to a strict either-or relation. This allows a participant to submit bids across multiple delivery windows at the same time. For instance, it can submit peak and off-peak blocks and have the algorithm select whichever maximises surplus, ensuring the asset is dispatched at the most profitable moment without the risk of double execution. It is for generators that do not have a preference of generation in a specific time frame, but rather aim for the market to decide for them the best time to sell an energy block.

Loop blocks are, at the moment, an order type exclusive to EPEX Spot which is designed to represent storage activities. Loop orders consist of pairs of block orders, one that buys and one that sells, that are executed or rejected together as a unit. One block represents the storage phase and the other the generation phase, allowing participants to bundle charging and discharging into a single conditional instrument. The order is only accepted if both legs clear simultaneously, ensuring that the round-trip economics are respected: a participant will not be dispatched to charge unless the corresponding discharge is also scheduled. Loop blocks enable storage participants (such as battery operators) to implicitly express their operating characteristics without explicitly disclosing them to the market operator, while accounting for the price uncertainty inherent in auction markets [9].

The Scalable Complex Order (SCO) is an order exclusively used by OMIE and EirGrid (Irish NEMO) [10][11]. It allows the definition of stepwise marginal cost curves per hour with a minimum acceptance volume per hour. SCOs include a fixed term and may be accepted even if they are out-of-the-money in some hours, provided that they are overall in-the-money over the delivery day when considering the fixed term and the hourly bid curves.

Category	Product / Condition	Technical Logic and Description
Smart Blocks	**Linked Blocks**	Establishes a parent–child hierarchy where a child block cannot be accepted if the parent is rejected. A profitable child can "save" an out-of-the-money parent if the total economic surplus is positive.
	Exclusive Groups	A set of blocks where the algorithm can execute only one (or a sum of acceptance ratios ≤ 1) to maximize social welfare.
	Loop Blocks	A family of two blocks (typically one buy and one sell) that must be executed or rejected together, designed for storage assets such as batteries.
	Scalable SCO	Stepwise marginal cost curves per hour with a minimum acceptance volume per hour. They also include a fixed term.

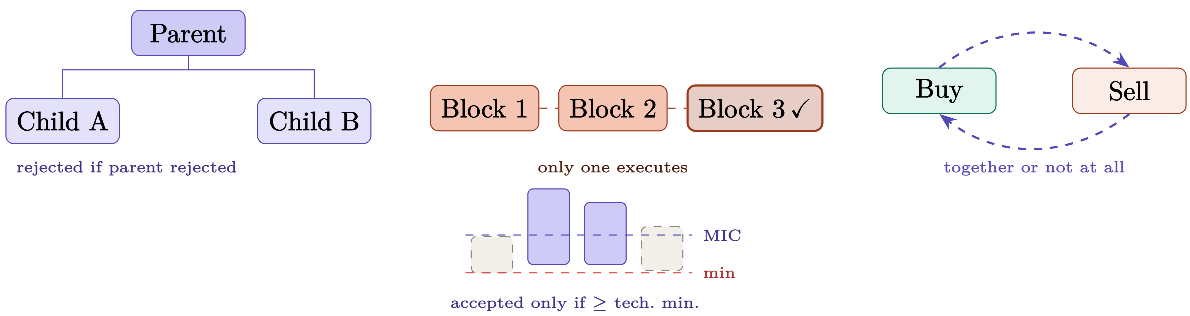

Linked (parent–child) · Exclusive (one of N) · Loop (buy+sell coupled) · SCO (stepwise curve with fixed cost floor)

2 | State of the art

This section examines Spain alongside four European markets: Germany, France, the Czech Republic, and Italy. These were chosen due to the fact that their corresponding NEMOs cover the main variants of auction order-type design currently in use across the continent. Adding to this, the selection of the countries themselves is not arbitrary: each country represents a distinct generation profile. Germany and Spain have large wind and solar shares; France depends on nuclear for nearly three quarters of its output; the Czech Republic runs primarily on coal and nuclear; Italy is the most gas-exposed market in Western Europe. These structural differences have shaped how each NEMO designed its product set, and comparing them places OMIE's order-type constraints in context.

The analysis proceeds in two steps. The first characterises the generation mix of each country, which determines the volatility and rebalancing pressure that the intraday market must absorb. The second examines the order types each NEMO offers in response. For the four non-Spanish markets, this characterisation relies on publicly available NEMO documentation; no unit-level offer data are accessible. Spain is treated differently throughout: OMIE publishes granular unit-level offer records, making it possible to state what Spanish generators actually submitted rather than what economic theory predicts.

2.1 | Data

The first analysis is a general one. The data was obtained from the ENTSOE Transparency Platform. This platform makes it possible to break down the generation mixes of each country for the whole year 2025 [12].

Table 2.1: Generation mix by country in 2025.

Technology	Spain	Germany	France	Czech Rep.	Italy
Renewables					
Solar	20.3	16.3	5.7	6.6	14
Wind	21.5	29.5	9.1	0.9	6.7
Geothermal	—	—	—	—	1.7
Other renewable	0.2	0.2	—	3.2	0.5
Nuclear					
Nuclear	20.0	—	69.6	42.4	—
Hydro (excl. pumped storage)					
Hydro	15.5	5.3	11.2	3.7	13.3
Fossil fuels + Waste					
Fossil gas	20.1	13.4	3.1	5.2	38.4
Fossil hard coal	0.5	22.0	—	33.3	4.3
Net imports	—	3.0	—	0.4	15.7
Biofuels	2.1	9.2	1.1	4.4	4.9
Oil & Other (non-renewable)	—	1.0	—	—	—
Total net generation	**100%**	**100%**	**100%**	**100%**	**100%**
Pumped hydro storage (PHS) — reference, excluded from total					
Hydro pumped storage (PHS)	3.5	2.7	1.7	1.4	1.1

Notes: (—) no data available, not applicable, or not broken out for that country in the source used.

This table already sets up some general conclusions.

Germany is the most fossil-fuel-dependent of the five despite its renewables build-out, with hard coal and gas together accounting for around 36%, reflecting the post-nuclear phase-out. Italy mirrors this with fossil gas alone at 38.4%, making it the most gas-exposed market. Gas also accounts for 20.1% of Spain's generation.

Hydro is meaningful in France, Italy, and Spain, adding a dispatchable renewable layer that interacts with intraday pricing. Pumped hydro, excluded from the totals, is present across all five but modest — Spain's 3.5% is the largest, relevant for BESS analysis since pumped hydro is the incumbent storage technology these markets were implicitly designed around.

The generation mix data in Table 2.1 contextualises the market design choices examined in the sections that follow; each country's order-type set reflects both the physical characteristics of its generation portfolio and the regulatory framework inherited from pre-liberalisation arrangements.

The Spanish electricity system

The first country which will be analysed is Spain. The Spanish electricity system has undergone a structural transformation over the past decade that places it among the most renewable-intensive markets in Europe. Solar photovoltaic and onshore wind together account for roughly 43% of net generation in 2025, a share that continues to grow as new capacity enters service each year. This penetration level is not uniform across the calendar: during spring and summer, solar generation can satisfy the majority of daytime demand, pushing residual load to near zero and producing frequent episodes of negative or near-zero prices in the day-ahead market. In winter, the balance shifts toward wind and reservoir hydro, the latter acting as a critical buffer when thermal generation is needed to cover demand peaks.

Nuclear provides a stable base-load layer of approximately 21%, operating as a near constant must-run technology that anchors the lower end of the supply curve throughout the year. Gas-fired combined cycle plants occupy the marginal position in the merit order for most peak hours, and their dispatch behaviour is closely tied to the spread between natural gas hub prices and day-ahead electricity prices. Reservoir hydro, at around 9% of annual generation, is perhaps the most strategically valuable asset in the Spanish system: operators manage it with a view to the entire seasonal water budget, deploying it preferentially in hours of high price and scarcity rather than as a continuous source. Reservoir hydro may also be dispatched according to its opportunity cost, close to the marginal price of gas when gas prices are high.

The Iberian wholesale market is operated by **OMIE** under the *Single Day-Ahead Coupling* (SDAC) framework, where it has been integrated since 2014. Spain and Portugal form a single bidding zone under normal conditions, though the interconnection between them is subject to congestion and can trigger a market split producing separate prices for each country. The matching algorithm for the countries mentioned, as explained before, is EUPHEMIA, which clears supply and demand on a marginal pricing basis for each of the 96 quarter-hourly periods now composing the trading day following the transition to MTU15 in late 2025.

In Spain, the intraday market averages around 18% of the day-ahead market volume. Within this segment, intraday auctions (IDAs) dominate, representing approximately 15%, while the continuous intraday market contributes only about 3%. This distribution reflects the structure of the market, which was substantially reformed in June 2024. At that point, the six legacy MIBEL regional auction sessions were replaced by three European IDA auctions which were integrated into the pan-European Single Intraday Coupling (SIDC) framework. Alongside these auctions, a continuous cross-border market operates via the XBID shared order book. This allows participants to adjust their positions up to one hour before physical delivery. Spain is one of the countries with the largest share of IDAs compared to continuous trading. This highlights the importance of auction-based price formation for the adjustments of the system.

OMIE is notable for using units of offer. It is also the only NEMO offering *complex orders in its day-ahead market* — a feature practically unique to OMIE — which include the Scalable Complex Order (SCO) and load-gradient constraints, instruments particularly relevant for combined-cycle operators recovering start-up costs and enforcing technical minimum generation levels. In its intraday auctions, however, OMIE offers only simple bids and simple blocks, in line with the product scope currently defined for SIDC Intraday Auctions across all participating NEMOs.

NEMO	Granularity	Simple bids	Complex bids
OMIE	15 min	Up to 25 segments per period	SCOs (Scalable Complex Orders) Simple blocks Exclusive groups

Table 2.2: Day-ahead market products (OMIE).

NEMO	Timings (IDAs)	Bid types (IDAs)
OMIE	15 min	Simple bids Simple blocks

Table 2.3: Intraday market products (OMIE).

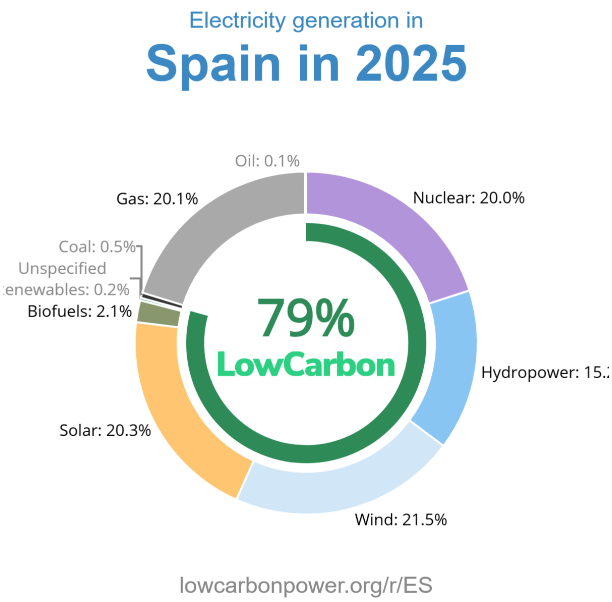

NEMOs operating in Germany: EPEX/Nordpool

Germany presents the most structurally complex electricity market in Europe, shaped by the pursuit of decarbonisation and security of supply following the definitive closure of its last nuclear plants in April 2023 and the increase in gas prices after the Ukraine conflict with Russia [13]. The generation mix is dominated by variable renewables; onshore wind at roughly 23%, solar at 16%, and offshore wind at 6%, alongside a thermal backbone of lignite (17%), natural gas (16%), and hard coal (7%) that continues to provide the dispatchable capacity that renewable sources cannot provide at the moment. Biomass contributes a further 9% in a largely scheduled, must-run role.

This mix shapes the market's price dynamics. On days of high wind and solar output, supply exceeds demand for most of the day, pushing marginal prices to zero or below; surplus generation that cannot be absorbed domestically must be exported or curtailed. On low-wind winter evenings, the same system faces scarcity, with gas and lignite setting prices well above the annual average. The resulting intraday price swings, driven by both solar-induced midday oversupply and multi-day wind variability, are among the main reasons Germany's intraday market has evolved into the deepest and most liquid in Europe.

Under the SDAC framework, EPEX SPOT operates the German day-ahead auction within a single bidding zone covering Germany and Luxembourg. Participants can submit simple hourly orders and block orders spanning multiple periods, with maximum block volumes of 600MW in the DE/LU zone. In addition, EPEX SPOT offers several complex block-type products in its day-ahead auction, including simple blocks, big blocks (larger than simple), linked blocks, exclusive blocks, loop blocks, and exclusive linked block families [14]. Day-ahead liquidity in 2025 remained stable despite renewable build-out: traded volumes reached 291,406.7 GWh, essentially unchanged from 291,403.2 GWh in 2024.

The intraday market is where Germany differs most: the SIDC/XBID continuous platform supports trading in hourly, 30-minute, and 15-minute contracts on EPEX SPOT, with trading typically possible up to a few minutes before physical delivery (five minutes for many bidding zones), enabling very short lead times by European standards. On the intraday auction side, coupled since 2024 under SIDC, the product scope is limited to simple bids and simple blocks, the same baseline currently applied across all SIDC IDA sessions. Intraday liquidity nonetheless kept growing in 2025: continuous intraday trading rose by 3.3% to 94,354.0GWh, up from 91,296.1GWh in 2024, while intraday auction liquidity increased by 9.7% to 11,652.2GWh. Combined, these developments pushed Germany's total intraday volume to a record 106,006.1GWh, compared with 101,915.6GWh the year before. This reflects the system's growing reliance on rapid re-adjustments as renewable penetration deepens [15] [16] [17].

NEMO	Granularity	Simple bids	Complex bids
EPEX / Nordpool	15, 30, 60 min	Linear price curves (linear or stepwise)	Smart blocks Linked blocks Exclusive blocks Flexible blocks Loop blocks (EPEX)

Table 2.4: Day-ahead market products (EPEX/Nordpool).

NEMO	Timings (IDAs)	Bid types (IDAs)
EPEX / Nordpool	15, 30, 60 min	Simple bids Simple blocks

Table 2.5: Intraday market products (EPEX/Nordpool).

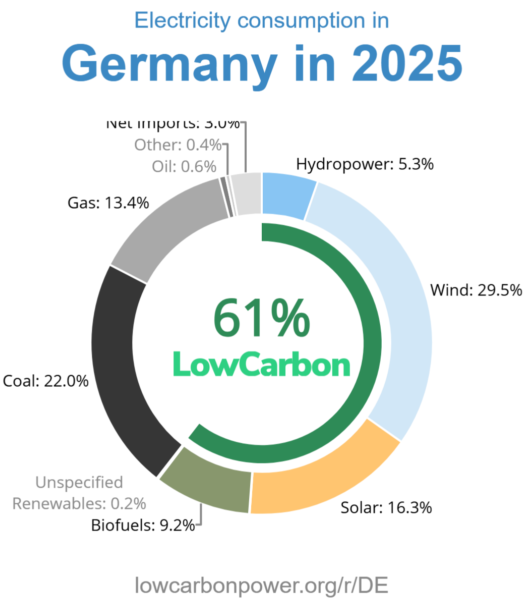

NEMOs operating in France:EPEX/Nordpool

France derives 69.6% of electricity from nuclear. It is the most nuclear-dependent grid in the world (The US has more nuclear power, but it doesn't cover such a high percentage of the mix). This is due to a legacy of strategic choices made in the decades following the 1973 oil crisis. This choice has produced a market whose price dynamics are fundamentally different from those of its neighbours: French baseload prices are systematically lower than the European average in most hours, reflecting the near-zero marginal cost of nuclear production, while winter peak prices can spike dramatically when demand surges coincide with reactor unavailability, as occurred during the exceptional maintenance crisis of 2022.

Gas, at roughly 3% of annual generation, is not a baseload technology in France but a peak and emergency resource. Hydro (combining reservoir and run-of-river assets at approximately 10%) is operated with considerable strategic discipline, its reservoir component withheld for hours of genuine price signal rather than deployed continuously. Wind, growing steadily toward 10% of generation, introduces a layer of intraday variability that the nuclear-heavy fleet is structurally ill-suited to absorb alone.

The French day-ahead market operates through EPEX SPOT as the designated NEMO, clearing within the SDAC framework at 12:00 CET. Beyond standard hourly energy bidding, EPEX SPOT offers a range of complex order types in day-ahead trading, such as blocks, big blocks, Smart Blocks, linked and exclusive blocks, and the newer "Exclusive Links" block families. All these are meant to allow market participants to represent technical constraints and portfolio strategies more precisely in the day-ahead auction. In 2025, however, day-ahead auction liquidity softened: traded volumes fell to 137,177.7GWh, down from 139,921.0GWh in 2024, reflecting a modest contraction in auction-based activity even as the broader market structure remained stable. Capacity in France is also subject to a capacity-obligation regime managed by RTE, with a secondary capacity market hosted by EPEX SPOT where suppliers must hold enough capacity certificates to cover their customers' peak demand, thereby providing an additional long-term investment signal alongside the energy-only day-ahead price [18].

The intraday structure mirrors Germany's: SIDC/XBID continuous trading in 15', 30', 60' granularity to minutes before delivery. Continuous intraday trading experienced a decline in 2025, with volumes falling to 15,735.0GWh from 17,094.7GWh the year before. However, this was due to a dramatic surge in intraday auction activity, due to their inclusion since end of 2024. Their liquidity rose by 150.8%, reaching 6,908.9GWh compared to 2,748.3GWh in 2024.

This explosive growth pushed the total French intraday market volume to 22,644.0GWh, up from 19,849.9GWh, showing a shift toward auction-driven balancing as renewable variability increases. As with the other SIDC participants, France's intraday auctions currently accept only simple bids and simple blocks; the richer complex order set described above for day-ahead trading is not part of the IDA product scope. The FR/DE cross-border interconnection distinguishes the French market architecture: contrary to the implicit capacity allocation used on most XBID borders, capacity between France and Germany is allocated explicitly, a regulatory choice which reflects divergent price levels and a need for tighter control over cross-border flows [14] [15].

NEMO	Granularity	Simple bids	Complex bids
EPEX / NordPool	15, 30, 60 min	Linear price curves (linear or stepwise)	Smart blocks Linked blocks Exclusive blocks Flexible blocks Loop blocks (EPEX)

Table 2.6: Day-ahead market products (EPEX/Nordpool).

NEMO	Timings (IDAs)	Bid types (IDAs)
EPEX / Nordpool	15, 30, 60 min	Simple bids Simple blocks

Table 2.7: Intraday market products (EPEX/Nordpool).

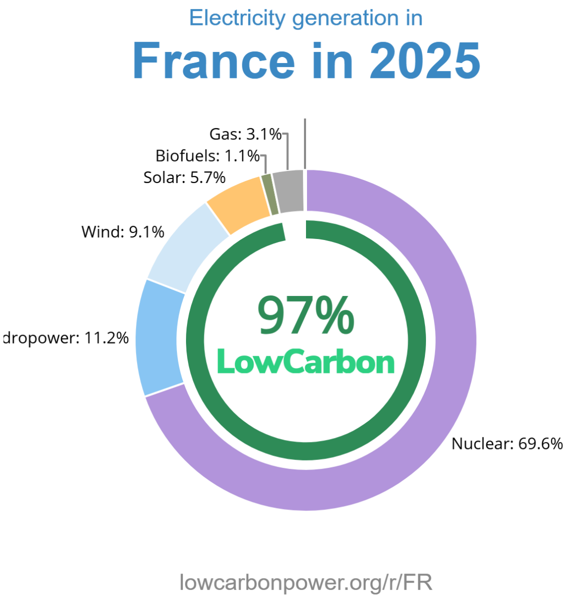

Czech Republic-OTE

The Czech electricity system is one of the most thermal and nuclear concentrated in Central Europe. Its generation profile reflects both the industrial legacy of the communist era and a deliberate post-transition policy of maintaining energy self-sufficiency. Nuclear plants supply approximately 43% of net generation, while lignite and brown coal account for a further 33% . Together, these two technology groups cover more than three-quarters of annual output with technologies characterized by very low marginal costs, high technical minima, and limited ramp capability. The consequence is a supply curve that is unusually flat and inflexible by European standards, with very little of the thermal capacity able to respond dynamically to intraday price signals.

Renewable energy penetration remains limited. Solar contributes roughly 7% of generation, and onshore wind less than 1%, figures that stand in contrast to the renewable shares observed in Germany or Spain. The effect is that the Czech market has not yet experienced the deep intraday volatility that characterizes more renewable-intensive systems. A key constraint is that the Czech Republic sits at the intersection of German, Polish, Austrian, and Slovak grids, experiencing frequent unscheduled flows (loop flows/OTF) that cause congestion independent of local dispatch.

The day-ahead market operates under the **4M Market Coupling** framework, a regional arrangement covering the Czech Republic, Hungary, Slovakia, and Romania, with price coupling services provided by EPEX SPOT in coordination with the national market operator **OTE**. Intraday participation through the SIDC continuous platform exists but remains less liquid than in Western European markets, partly because the inflexible generation base creates limited commercial incentive for active intraday repositioning.

NEMO	Granularity	Simple bids	Complex bids
OTE	15 min	Step orders	Profile block bids (PBO) Exclusive groups Linked blocks

Table 2.8: Day-ahead market products (OTE).

NEMO	Timings (IDAs)	Bid types (IDAs)
OTE	15 min	Simple bids Simple blocks

Table 2.9: Intraday market products (OTE).

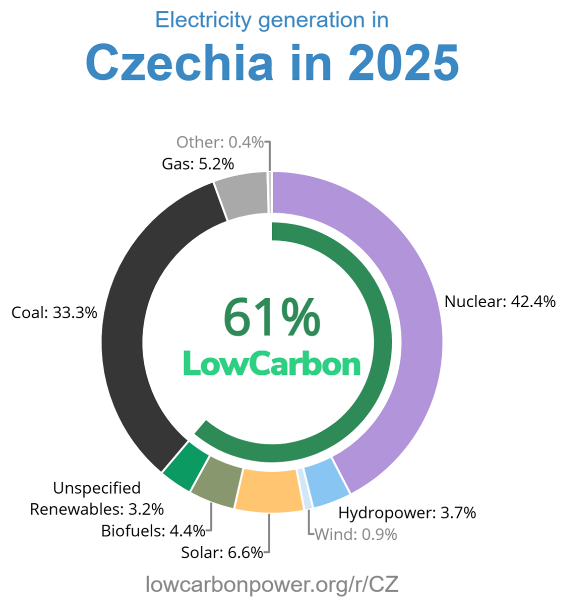

Italy-GME

Italy is the only major Western European economy that operates a large interconnected grid without any nuclear generation. It reached this position through two referendums, in 1987 and again in 2011, that effectively closed its nuclear plants. The structural consequence is a generation system highly dependent on natural gas, which supplies approximately 38-39% of net electricity in 2025. Gas is not merely the marginal technology in Italy: it is the price-setter in the vast majority of hours throughout the year, meaning that Italian wholesale electricity prices are more tightly correlated with the TTF gas hub than those of any other major European market [19] [20].

This gas dependence is partially offset by a diverse hydro base: run-of-river plants contribute roughly 12% of annual generation from the alpine watersheds of the north; and by a growing solar sector now approaching 16% of net output. Onshore wind, concentrated in the centre and south of the country, adds a further 10%. The Italian renewable profile is spatially heterogeneous. It is solar-heavy in the south and centre, hydro-heavy in the north, with generation patterns that do not necessarily align with the distribution of industrial demand concentrated in the Po Valley. This mismatch is one of the drivers of the structural transmission congestion that defines the Italian market's most distinctive institutional feature: the division of the country into seven separate price zones [21].

The zonal structure, operated by **GME** under the *Mercato del Giorno Prima* (MGP), means that a single clearing price does not exist in Italy in the way it does in Spain or France. Instead, each zone clears at a price reflecting the local balance of generation and demand, constrained by transmission capacity between zones. The PUN (national reference price) was discontinued in 2023; zonal prices now determine all revenues and costs [22] [23].

In Italy, the intraday market (MI) (*Mercato Infragiornaliero*), operated by Gestore dei Mercati Energetici, accounts for a relatively limited share compared to other European countries like Spain or Germany, with total intraday volumes reaching 35,400 GWh in 2024 (at the moment of writing this thesis, no data for 2025 regarding this subject has been published). This corresponds to approximately 16% of day-ahead traded volumes (226,800 GWh). Within the intraday segment, auction-based trading (IDAs) remains dominant, accounting for 23,900 GWh, or around 10.5% of day-ahead volumes, while continuous trading via XBID reached 11,500 GWh, representing approximately 5% of the day-ahead market. This distribution reflects the structure of the Italian intraday market, which was significantly reformed in June 2024, when the national MI-A (CRIDAs) auction sessions were integrated into the European Intraday Auctions (IDAs) under the Single Intraday Coupling (SIDC) framework. Alongside these auctions, continuous cross-border trading takes place via the XBID shared order book, enabling market participants to adjust their positions closer to real time. The stronger weight of auction volumes compared to continuous trading highlights the continued importance of structured auction mechanisms in Italy, although the rapid growth of XBID points to increasing liquidity and participation, particularly from renewable energy sources.

It is also worth noting the TIDE implementations from February 2026, which coincided with the development of this thesis and led to substantial changes regarding GME. It addressed the imbalance settlement period (ISP), which now supports finer granularities: time intervals for bids/offers (for both simple and block products) have been set to 15'-30'-60' for the MGP, 15' for the MI-A and 15'-30' for the XBID. A second introduction of the TIDE is the refined block bid/offer methodology. GME now allocates maximum block bids/offers quarterly for both MGP and MI-A sessions; a base share goes evenly to active participants (those submitting valid bids in the prior observation period) with remaining blocks distributed proportionally to the number of portfolios held. Weights reward prior activity (e.g., 1.0x base for participants using ≥50% of prior blocks in any session, 0.7x otherwise) [24].

NEMO	Granularity	Simple bids	Complex bids
GME	15, 30, 60 min	Maximum 4 bids per interval and portfolio	Simple profile blocks of 15, 30 and 60 min

Table 2.10: Day-ahead market products (GME).

NEMO	Timings (IDAs)	Bid types (IDAs)
GME	15 min	Simple bids Simple blocks

Table 2.11: Intraday market products (GME).

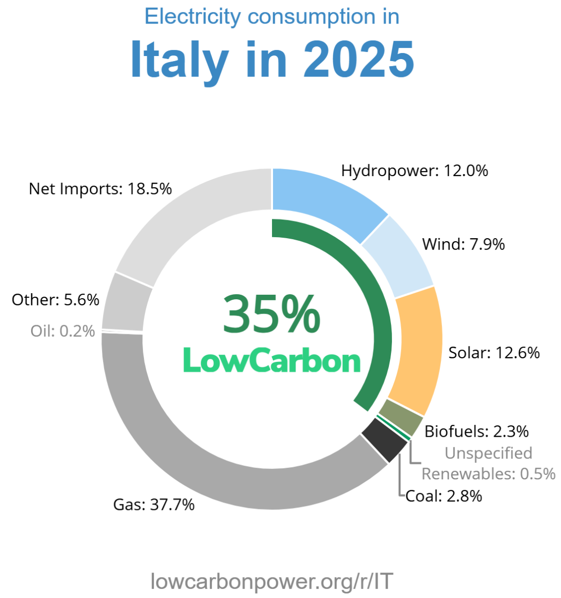

2.2 | Bidding behaviour and market design implications

The bidding behaviour across all five markets follows the standard economic framework for generator offer strategies in deregulated electricity markets. Technologies with negligible variable costs bid at or near zero to guarantee dispatch; thermal generators offer at their short-run marginal cost, primarily fuel and CO_2 costs [25]; reservoir hydro offers at its opportunity cost of water rather than at operating cost [26]. Where the market structure permits, thermal plants with significant start-up costs also use complementary order types to recover fixed start-up costs [27]. The OMIE offer data can confirm this directly for Spain's day-ahead market: Table 2.12 shows wind and solar submitting at approximately zero, nuclear operating as must-run with SCO-backed minimum income guarantees, and CCGTs as the dominant users of all three complex order types. The theoretical expectation that variable-cost technologies bid near zero [28] is therefore not just a modelling assumption but an empirically observed pattern in the one market where unit-level data is available. Only Day-Ahead and Intraday Auction bidding is considered; the continuous intraday market falls outside the scope of this thesis.

For Germany, France, the Czech Republic, and Italy, no empirical offer data were available: none of the corresponding NEMOs (EPEX SPOT, OTE, or GME) publish public unit-level offer records. The descriptions of those four markets therefore draw on the economic framework above, not direct observation.

Spain is treated differently, as OMIE publishes granular unit-level offer data through its public data portal (in "resultados del mercado"-> "acceso a ficheros"), including the DET and CAB settlement files that record the specifics of the offers sent by the agents. OMIE also publishes aggregate offer-typology statistics covering different periods [29]. The characterisation in Table 2.12 is derived directly from this data: it states what Spanish generators actually did, not what they are expected to do.

Table 2.12: Observed bidding behaviour — Spain (OMIE), March 2025 – January 2026

Market structure: unit-level offers — technology attribution is directly observable from DET and CAB files.
DA order types available: Simple bids (25 segments), SCOs, Simple blocks, Exclusive groups
Intraday order types available: Simple bids, Simple blocks (IDA)

Technology	**Mix**	**Day-ahead bidding behaviour**	**Intraday behaviour**
Onshore wind	21.5%	Simple offer at ≈0 €/MWh; not observed in any complex order category	Simple bids in IDA
Solar PV	20.3%	Simple offer at ≈0 or negative; not observed in any complex order category	Simple bids in IDA
Nuclear	20.0%	Simple offer at ≈0 (must-run); `SCO` observed consistently (≈2–3 offers/day) for minimum income guarantee (FijoEuro + load gradient) [27]	Minimal intraday
Gas (CCGT)	20.1%	Simple offer at marginal cost; dominant user of all three complex order types: `SCO` (≈28–35 offers/day), `Exclusive Group Block Order` (≈15–25 offers/day), and `Simple Block Order` (variable); consistent with start-up cost recovery and multi-hour commitment needs [27]	Simple bids and simple blocks in IDA
Reservoir hydro	∼9%	Simple offer at opportunity-cost price; `simple blocks` and `exclusive groups` available for production-window and unit-level constraints; limited presence in complex order data	Simple bids in IDA
Pumped hydro	3.5%	Simple bid/offer; `simple blocks` available for charge/discharge window coordination; limited presence in complex order data	Simple bids and possibly simple blocks in IDA
Biomass	2.1%	Simple offer at low stable price; occasional `Simple Block Order` observed	Minimal intraday
Hard coal	0.5%	Simple offer; `Exclusive Group Block Orders` and `Simple Block Orders` observed; residual capacity	Very limited intraday

The order types available to each of these technologies reflect the assets the market was designed around. However, this is something that has been studied for the past years and which can be assumed to have a certain academic background. There is little to no empirical research conducted on the complex orders sent in IDAs (although they only include blocks). The fact about which technologies and how they decide to reposition themselves in the intraday market is less researched than the day-ahead market. Although they face the same technological constraints, they might decide to operate differently. And for newer assets, such as battery storage systems, these sets of constraints are different. Battery storage is now becoming commercially active across all five markets. Deployed capacity has increased substantially throughout Europe (especially in Germany, France, Italy, and the United Kingdom) over the past three years [30], and the Clean Energy Package redefined storage as a distinct market category for the first time a few years ago [31].

Across the five markets, day-ahead product sets differ substantially: EPEX SPOT offers the richest range, including linked blocks, smart blocks, exclusive blocks, flexible blocks, and the Loop Order (C88); OTE and GME offer simpler products; OMIE relies on its SCO and exclusive-group instruments. In intraday auctions, however, this variation largely disappears: under the current SIDC IDA framework, all five NEMOs — including EPEX SPOT — offer only simple bids and simple blocks, a scope set at the European regulatory level rather than by individual NEMO choice. OMIE's IDA product set is therefore not restrictive relative to its peers; it matches the common baseline currently in force. Further on it will be evaluated whether this common baseline is adequate for participants with intertemporal operating constraints, using the Spanish case, where unit-level data make empirical analysis tractable. Chapter 4 exploits this to characterise OMIE's intraday market in detail: which technologies participate in each session, how frequently block orders are used, by whom, and for what operating reasons. That empirical layer underlies the central question this thesis addresses: whether the simple-bid/simple-block scope common to SIDC intraday auctions is adequate for participants with intertemporal operating constraints, and what a more expressive instrument, comparable to the Loop Order available in EPEX SPOT's day-ahead auction, would contribute to market efficiency if introduced at the intraday level. The thesis uses a battery model as the analytical example: a storage asset whose intertemporal charge-discharge structure might not be able (at the moment) to be fully expressed through the order types currently available in intraday auctions. This will provide a clean test case for where the product design falls short.

The analysis will thus require two inputs: a price forecast to recreate how a battery operator would operate in the intraday auctions without the certainty of the prices that will appear, and an optimisation framework built upon these order types. The following two subsections review the literature underpinning each.

2.3 | Electricity Price Forecasting

Participation in OMIE's intraday auctions requires submitting bids before session closure, when clearing prices are not yet known. The price forecast therefore enters every decision in the optimisation model: it determines whether the battery bids at all, which order type is selected, and at what price level. The following review situates the methodological choice made in Section 4.5.

The field of electricity price forecasting (EPF) has produced three broad families of methods since Weron's review [32]. Statistical methods, autoregressive models (AR, ARIMA) and their regularised extensions, use historical prices and calendar features to project price dynamics through a linear structure. Machine learning methods, principally deep neural networks (DNN), long short-term memory networks (LSTM), and convolutional networks (CNN), learn non-linear mappings without imposing parametric form. Hybrid methods combine decomposition algorithms, feature selection, and one or more forecasting models, typically at the cost of interpretability.

Lago's 2021 forecasting benchmark provides the most rigorous cross-market comparison of these families to date, evaluating methods on five datasets each spanning six years and reporting both point-forecast accuracy and statistical significance tests [33]. Their central finding is that the Lasso Estimated AutoRegressive (LEAR) model — a parameter-rich autoregressive specification estimated with ℓ_1 regularisation, originally proposed by [34] — matches DNN point-forecast accuracy on day-ahead European markets while, in their benchmark setup, requiring 30 to 100 times less computation. LEAR fits one regularised regression per time period rather than a single model across all periods, so the price dynamics of the 08:00 slot are estimated independently from those of the overnight hours. The ℓ_1 penalty automatically eliminates redundant regressors; the estimator is therefore robust to overfitting even on short training windows, a practical advantage given that the available intraday sample is limited to the period after OMIE's IDA sessions began. Lago notes that DNN ensemble methods achieve statistically superior accuracy overall, but that LEAR is the dominant choice when decision-making must occur within seconds [33] which is the case for bid submission ahead of intraday auction gate closure.

The per-period LASSO framework, established for day-ahead forecasting [35], has since been applied successfully to intraday markets [36], supporting its suitability for the present application.

This thesis adopts LEAR as the price forecasting layer following the specification in [33], with daily retraining on an expanding window. The feature set follows the standard 17-regressor LEAR specification and is extended with the same-period day-ahead clearing price as an 18th feature, known before any IDA session opens, which anchors the intraday price level. Forecast accuracy is reported as mean absolute error (MAE) over the walk-forward evaluation period and discussed in Chapter 5.

2.4 | Battery Storage Self-Scheduling and Bidding

The optimisation framework that translates price forecasts into bidding decisions draws on two strands of literature: deterministic self-scheduling under perfect foresight, which establishes the upper bound on extractable value, and stochastic bidding under price uncertainty, which describes the realistically implementable problem.

The deterministic self-scheduling problem for a price-taking storage asset is formulated as a linear or mixed-integer linear programme (LP/MILP) over a finite time horizon. The objective is to maximise arbitrage revenue by charging when prices are low and discharging when prices are high while being subject to state-of-charge dynamics, power limits, and round-trip efficiency losses. Sioshansi et al (2009) formalise this LP structure for a price-taking battery operating in a day-ahead pool and show that the optimal policy takes a threshold form: the battery charges whenever the price falls below a buy threshold and discharges whenever it exceeds a sell threshold, where both thresholds are jointly determined by round-trip efficiency and any operating cost [37]. This structural result motivates the welfare benchmark W^* in Section 4.6, which solves the LP under perfect foresight against the realised IDA prices to obtain the maximum energy value accessible without any bidding constraint.

The stochastic extension consists of submitting a bid before prices are known. This requires the strategy to be optimal in expectation across possible price realisations rather than for the ex-post realised price. Scenario-based stochastic programming approaches for electricity market participants, including storage assets, in which a finite set of price trajectories drawn from a forecast distribution is used to evaluate the expected revenue of each candidate strategy [38]. A key result is that realised revenue falls short of W^* for two separable reasons: forecast error, because the submitted bid is optimised against the forecast rather than the realised price; and order-type constraints, because the available instruments may not be expressive enough to implement the LP-optimal dispatch even under perfect foresight. Separating these two contributions requires W^* as a common upper bound and a methodology that evaluates each strategy ex post against realised prices, which is the approach adopted in Section 4.6.6.

Existing storage bidding models in the literature typically assume that the asset can submit an arbitrary price-quantity schedule independently for each period, equivalent to a set of simple orders with no intertemporal coupling [37] [38]. This assumption is reasonable in continuous intraday markets where bids can be updated close to physical delivery, but it breaks down in single gate-closure auction sessions. A storage asset that charges in one period is physically committed to discharge the stored energy later; simple orders submitted independently for charge and discharge periods could possibly not be enough to express this joint constraint. If only one leg is matched, the asset faces execution risk. In EPEX SPOT's day-ahead auction, the Loop Order (C88) addresses this directly by linking a buy block and a sell block into a single accepted-or-rejected unit, ensuring that neither leg is accepted without the other [14].

No equivalent instrument currently exists in any SIDC intraday auction. Notably, day-ahead auctions (SDAC) and intraday auctions (SIDC) are cleared by the same EUPHEMIA algorithm, which already supports these complex order types on the day-ahead side; extending their availability to intraday auctions would therefore be an extension of an existing algorithmic capability rather than the development of a new one. Whether introducing one into OMIE's IDA product set would improve market efficiency for participants with intertemporal operating constraints is the central empirical question this thesis addresses through the welfare gap analysis in Chapter 5.

3 | Problem setting

3.1 | Intraday Market Context and Research Gap

The existing literature on intraday auction design is largely theoretical: it describes the order types available across NEMOs and the economic logic behind them, but offers little empirical evidence on how these instruments are actually used in practice. Of the five markets examined in Chapter 2, only OMIE publishes unit-level offer data, which means that for four of the five countries the question of which technologies use which order types, and how successfully, has no publicly verifiable answer. However, it is possible to evaluate whether any of their options can be of use for OMIE to implement. This empirical gap is the starting point for the analysis that follows. However, in this thesis, a closer focus will be pointed towards intraday auctions. With the increasing penetration of variable renewable energy sources, as well as the necessity of introducing other generation types when the renewables encounter unexpected scenarios (heavy cloud cover, the sun/wind not being as strong as forecasted), these auctions are becoming more and more critical to adjust these positions.

In high-renewable penetration systems such as Spain and Germany, a significant share of generation is initially scheduled in the day-ahead market and subsequently adjusted in intraday auctions. The analysis will focus in detail on the typology of IDAs: what each instrument allows a participant to express, where its constraints lie, and how well it maps onto the operating characteristics of different technologies. The technologies that use each order type, and the extent to which their intraday volumes relate to their day-ahead positions, are examined as evidence of how well the existing product set serves them.

The order types offered in IDAs are not equally suited to all participants. Simple bids, submitted independently for each period, are adequate for technologies whose output in one period carries no operational consequence for the next — this is the case for wind and solar, which are expected to use intraday auctions mainly to correct forecast deviations period by period. Simple block orders, by contrast, exist precisely because some technologies might opt to be represented differently than always by these independent per-period bids: a block lets a participant commit to a single price across several consecutive periods, which matters for assets with ramping constraints, minimum-run times, or costs that are only recovered across a multi-period block rather than within a single period. Conventional generation and other flexible assets are the technologies expected to rely most on this second instrument, since their operating constraints span multiple periods rather than resetting each quarter-hour.

Whether this two-instrument set is actually sufficient in practice is the open question. The market design offers simple bids and simple blocks, but it remains unclear how each is used: which participants rely on blocks rather than simple bids, how often submitted orders of each type actually clear, and whether the instruments available are expressive enough to represent the intertemporal constraints these participants face. Existing literature provides limited empirical evidence on whether the available order types are effectively aligned with the operational needs of different technologies. This creates an important gap that requires systematic analysis.

The core tension that arises from these observations is that, while intraday auctions are critical for renewable-heavy grids (expected to operate with a simpler typology) the order types they make available might need to be revisited, to keep pace with what participants actually need.

3.2 | From Observed Market Behavior to Missing Instruments

Building on the empirical analysis of IDAs, a central question emerges: are the currently available order types sufficient to represent the operational and economic needs of market participants?

In this context, the increasing relevance of energy storage, and especially battery energy storage systems (BESS), reinforces the need to reassess intraday market design. Spain's National Integrated Energy and Climate Plan 2023–2030 (PNIEC 2023–2030), updated by Royal Decree in 2024, sets a target of 22.5 GW of installed energy-storage capacity (including batteries, pumped hydro, and solar-thermal storage) by 2030, as summarized in the EU-wide NECP 2025 compilation and related official communications [39] [40]. Complementing these policy targets, industry projections indicate a rapid acceleration of battery deployment: according to APPA Renovables, battery storage tied to self-consumption grew 119% year-on-year in 2025 (155→339 MWh), signalling rapid uptake at the distributed scale even as utility-scale deployment remains nascent. This shows that there will need to be increasing policy support, but also market-driven investment signals [41]. Similar trends can be observed across Europe, where storage is increasingly recognized as a key enabler of renewable integration and system flexibility.

Despite this expected growth, OMIE's current market design does not yet provide bidding instruments specifically adapted to the operational characteristics of storage technologies. In both day-ahead and intraday markets, batteries must rely primarily on simple orders, typically arbitraging price differences by charging during low-price periods (e.g., night hours or solar valleys) and discharging during peak demand periods. While this strategy captures a large amount of the available value, it does not allow storage operators to express the intertemporal constraints and minimum revenue requirements that define their viable charge–discharge cycles. As a result, there is a mismatch between the increasing system relevance of storage and the limited expressiveness of available market instruments.

The urgency of this gap is not merely theoretical. Industry analysis from consultancies such as AFRY projects that 2026 will mark a year of high-scale deployment in Spain, with installed capacity expected to reach 500–700 MW by year-end, rising to over 1 GW in 2027. This acceleration is driven by two converging trends: falling LCOS (levelized cost of storage) due to declining battery capex and extended technical lifetimes, and widening intraday price spreads produced by solar cannibalization. The resulting arbitrage opportunity is now sufficient to justify merchant investment even without revenue guarantees, yet the market structure through which these assets must bid has not kept pace. As Revuelta & Cobo observe, "a battery should not be understood as such [as selling solar energy during peak hours]; the sales revenue should conceptually be assigned to the margin that the battery will generate". In other words, a battery's revenue model is fundamentally a spread-arbitrage across coupled periods, not a simple generation sale, and the bidding instrument should reflect that structure [42].

In OMIE's intraday auctions, market participants primarily rely on simple and simple block orders, which allow independent bidding in each time period. While this structure provides flexibility, market participants might wish to express intertemporal constraints that might arise or dependencies across periods. This limitation becomes particularly relevant for assets whose profitability depends on coordinated decisions across multiple time periods. Among these, battery energy storage systems (BESS) provide a clear example: their value depends on executing consistent charge–discharge cycles, which might not be fully captured through independent hourly bids. As a result, simple orders can potentially expose such assets to risks including partial execution ("half-cycles") and suboptimal dispatch.

EPEX SPOT and *Nord Pool* offer other order types in their market designs, including linked blocks, or the variation of linked blocks for tying buy and sell offers known as loop blocks, which allow participants to express intertemporal constraints and a conditional execution. However, there is not yet any dedicated or adapted order type for storage assets currently existing in OMIE's (neither day-ahead nor) intraday auctions, despite the increasing relevance of such technologies in the energy transition.

This creates a second research gap: the absence of order types capable of representing multi-period operational constraints and revenue requirements of storage systems in (either day-ahead and) intraday markets.

3.3 | System Description

The second part of the analysis focuses on a grid-connected BESS participating in OMIE's intraday market (specifically the intraday auctions known as IDAs) under a pay-as-cleared settlement mechanism. The battery is parameterized by its rated capacity, maximum charge and discharge power, round-trip efficiency, and state-of-charge bounds. The full parameter values are given in Section 4.6.1. The operator is assumed to be a price-taker: their bids and volumes submitted do not affect clearing prices.

The core modeling question is how to translate a physically derived charge–discharge schedule into feasible bidding strategies under different order-type regimes, and what revenue each type captures relative to the unconstrained physical optimum. The regimes compared include orders already available, such as Simple Orders, which place independent period-level offers; and Simple Blocks, which couple decisions across periods. The last order evaluated will consist of Loop Blocks, which are not available at the moment, and enforce joint acceptance of a full cycle. The welfare gap between each regime and the LP-optimal dispatch quantifies the revenue cost of the order structure.

3.4 | Research Questions

As ACER's methodology for market product design is constantly being updated, NEMOs must ensure that the products they offer remain aligned with market participants' actual and potential bidding needs. This implies a constant assessment of which order types and market products should be retained, modified, or discontinued, based on the use and necessities of market agents. Against this background, the thesis focuses on the intraday market and evaluates whether the current product design adequately reflects the needs of its participants, as well as whether additional or more specialized products may be justified.

The thesis is guided by a main research question with five subordinate questions that follow a clear logic.

Main RQ: What is the current state of intraday auction market design, and how should OMIE adapt their products to reflect actual and emerging bidding needs under evolving ACER methodology?

Sub-RQs:

1. What role do intraday auctions (IDAs) play in adjusting day-ahead positions, and which participants actively use them?
2. Which existing intraday auction products are actually used, and what participant needs do they serve?
3. Do OMIE's block order parameters serve all technologies equally well, or are clearance outcomes systematically explained by structural cost differences (fuel/start-up costs vs. storage arbitrage economics) across technology types?
4. Are simple orders and simple block orders enough to adequately reflect the needs of market participants in intraday auctions or, on the other hand, would more complex order types improve market design?
5. Given the expected increase of batteries and storage assets in the following years, would a specific order type for them be justified? The thesis evaluates this primarily in the intraday auction context, where the product gap is most acute, and validates the same framework against the day-ahead market as a methodological baseline.

4 | Analysis and Proposed Methods

This chapter will develop four components. The empirical block-order analyser (Section 4.4) characterises complex-order usage across one year of Spanish IDA auctions and identifies which asset classes use block orders and for what operating reasons. The second aspect, which ties directly to the battery operation model consists of a price forecasting layer (Section 4.5). This produces the intraday price trajectories that feed the optimisation models. The welfare benchmark (Section 4.6) computes the maximum arbitrage revenue a perfectly informed, unconstrained battery could earn; this is the upper bound against which all bidding strategies are measured. The stochastic bidding model (Section 4.8) selects the revenue-maximising order type day by day under price uncertainty; this includes a Loop Order (C88), a paired charge/discharge instrument available in EPEX SPOT but not currently in OMIE, evaluated as a policy counterfactual for bid expressiveness. Revenue earned under each order type and its gap from the welfare benchmark are the central quantities reported in the results chapter.

The combination of these components enables a structured comparison between observed market usage and model-driven operational strategies. The backcast reveals which order types are actively used and by which technologies, while the battery model tests whether those same instruments are sufficient to represent the operational logic of a storage asset. Any systematic limitation in aligning these two perspectives points to a potential gap between market design and technological capability. Identifying this gap motivates more expressive bidding instruments.

4.1 | Implementation of the Backcast Validation Framework

4.1.1 | Purpose and Role in the Thesis

The backcast validation framework is introduced to bridge the gap between formal market design and observed market practice. Starting from raw OMIE auction files, it reconstructs unit-level dispatch (prices submitted, quantities matched, and order types used) directly from the published clearing results. This reconstruction is only possible because OMIE releases a uniquely detailed set of public market data; in many other European markets operated by different NEMOs, equivalent information is either aggregated or kept confidential. The Spanish case therefore offers a rare opportunity to analyse market behaviour at the level of individual units and order types.

The framework serves two distinct needs.

The *first* need was methodological. It was the first backcast developed in this thesis, as the main goal was to confirm that the parsing and interpretation of OMIE files were correct at the level of precision required for the optimisation that follows. This validation was performed for both the day-ahead and intraday market on a single reference day (1 December 2025) using first the day-ahead auction files, whose format is simpler and more thoroughly documented than the intraday equivalents. The daily backcast (`build_backcast.py`) then reconstructs unit revenues using the formula in Section 4.3.4, verifies SCO and Exclusive Block compliance for every complex offer, and cross-checks the results against independent calculation for different generators. These checks were also performed against OMIE's own database (which includes more detail for each independent agent and is easier to access than the published files). The tool also functions as a general-purpose inspection layer: any participating unit can be queried by generation type or offer typology, and the dashboard exposes prices, quantities, and order-type classifications for any date in the sample.

The *second* need is substantive: to characterise how block orders are actually used in the intraday market, before any modelling choice is made. The intraday auctions differ materially from the day-ahead in their product space (exclusive groups and SCOs are absent, leaving the block order as the only complex-order instrument) and it is in these sessions that the thesis's optimisation model operates. By iterating over one year of ICAB and IDET files and aggregating every submitted block order by unit, technology, direction, session, and size, the yearly block-order analyser (`yearly_data.py`) establishes whether blocks are a marginal curiosity or a structurally important tool for assets with multi-period operating constraints.

The two tools therefore form a sequence: the day-ahead backcast validates the pipeline on a known, well-documented case; the yearly intraday analyser then applies the same parsing logic at scale to answer the empirical question that motivates the order-type comparison in the results chapter.

4.2 | Software Architecture

The design of both tools is most easily understood in the context of the OMIE files they use.

4.2.1 | OMIE file types

The backcasting model operates with six families of files published daily by OMIE through its public data portal. All files follow the naming convention `<TYPE>_YYYYMMDD[SS].1`, where `SS` is the two-digit intraday session index where applicable. For day-ahead files the session suffix is absent. Table 4.1 summarises the role of each family in the model.

Table 4.1: OMIE file families used in the backcasting model

File	Content	Role in model
PDBC	Day-ahead matched dispatch programme	Baseline dispatch commitment for each market unit after the DA auction clears; starting point for intraday simulation
marginalpdbc	Day-ahead clearing price (€/MWh per period)	Uniform settlement price for DA revenues; benchmark for intraday offer evaluation
CAB / ICAB	Offer header metadata (DA and intraday)	Identifies offer type (simple and complex), required average price for block orders, and minimum acceptance fraction
DET / IDET	Offer bid curves—price-quantity staircases (DA and intraday)	Primary input to clearing simulation; each row is one step of one offer in one period
marginalpibc	Intraday auction clearing price (€/MWh per period, per session)	Used to evaluate simple-offer clearing conditions period by period and block-order acceptance over the block's span
PHF	Final hourly programme after system-operator adjustments	Used exclusively for generation mix visualisations; excluded from clearing simulation as it incorporates non-market dispatch decisions (like technical restrictions imposed by Red Eléctrica

With this file taxonomy in place, the design of each tool follows naturally. Both tools, the daily backcast (`build_backcast.py`) and the yearly block-order analyser (`yearly_data.py`), are written in standard Python 3.9 with no external dependencies beyond the interpreter's standard library. Chart rendering is delegated to Chart.js, loaded at runtime from a public CDN, so the final HTML files are portable and do not require a local server.

The overall architecture follows a strict separation between data extraction (Python) and presentation (JavaScript embedded in a templated HTML string): Python reads the OMIE files, aggregates them into a single dictionary, serializes that dictionary to a JSON string, and substitutes it into a placeholder inside an HTML template. The browser receives a frozen snapshot of the period analysed and renders all interactive tabs, tables, and charts so any user can visualize it. This design avoids any dependency on pandas, NumPy, or a web framework, and keeps each artifact under a single `.py` file that can be shipped along with the thesis. Figure 4.1 summarises this pipeline.

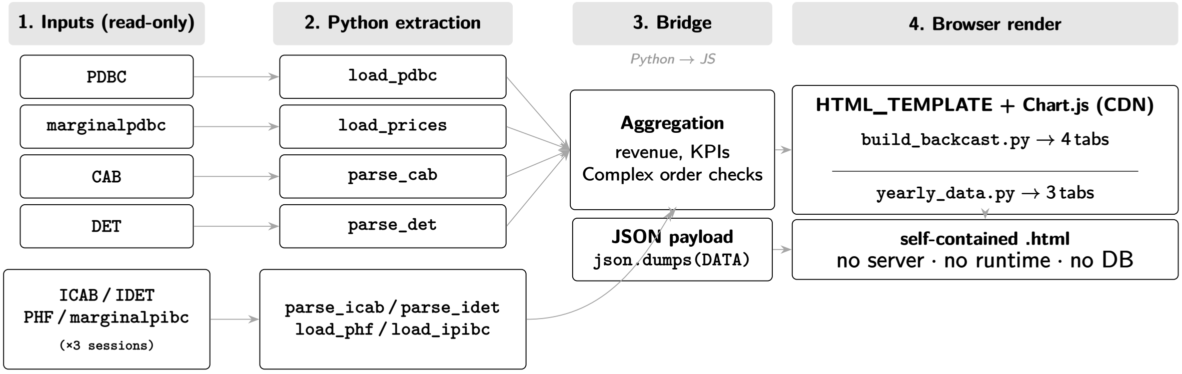

Figure 4.1: Software architecture shared by `build_backcast.py` and `yearly_data.py`.

4.3 | Daily Backcast: `build_backcast.py`

The daily backcast is structured as a linear pipeline: configuration, technology mapping, file ingestion, revenue computation, and dashboard rendering.

4.3.1 | Technology Mapping

Before any file is read, the script defines a large dictionary `TECH_MAP` that associates each OMIE unit code (the seven- or eight-character alphanumeric identifier such as `ALZ1` or `EGVD134`) known as the "Unidades de Oferta" with a technology label (*Nuclear*, *CCGT*, *Pumped Hydro*, *Wind*, *Solar*, *Retailer*, and related categories). This mapping was constructed by cross-referencing the public OMIE unit registry (found in an Excel file called LISTA_UNIDADES) with the internal classification of generation technologies used for the backcast.

4.3.2 | Day-Ahead Data Ingestion

Day-ahead ingestion is split across four parsing routines, each specialized to one of the OMIE file formats. The functions are deliberately small and defensive: they read line by line, strip trailing semicolons, skip headers and malformed rows, and silently ignore any line that cannot be parsed.

- **`load_pdbc`.** Reads the matched day-ahead program and returns a nested dictionary indexed first by unit code and then by period, with the signed matched megawatt value as the scalar entry. Positive values correspond to generation, negative values to consumption or pumping. The routine uses Python's `collections.defaultdict` to avoid key-existence checks downstream: any unit not matched in a given period is implicitly zero.
- **`load_prices`.** Reads the marginal-price file and returns a flat {period → price} dictionary. The Spanish clearing price is the fifth field of each data row; the routine enforces that the period index is numeric before attempting the float conversion on the price itself.
- **`parse_cab`.** The offer-header file is considerably less regular than PDBC or the price file because it mixes fixed-width fields, whitespace-separated tokens, and an offer-direction marker (`VO` for *venta*, a sell offer, or `CO` for *compra*, a buy offer). The parser locates this marker with a regular expression and extracts two numeric

tokens immediately after it: the declared minimum income of the complex offer (`ingreso_min`) and the committed energy (`e_casada`).

The unit code is recovered by scanning the portion of the line *before* the marker for any of the unit codes that appeared in PDBC. This cross-reference resolves the well-known difficulty that CAB uses an internal offer identifier (`cab_id`) which, by itself, is not sufficient to link back to a physical unit; the PDBC-based scan ensures that only offers belonging to units that actually participated in the matching are retained. The function returns two dictionaries: `cab_id_to_cod`, which links each offer identifier to its physical unit, and `cab_data`, which stores the minimum-income threshold and matched energy per unit. This is the source to find complex offer specs, such as the minimum-income condition used throughout the SCO compliance tab.

- **`parse_det`.** The detail file contains the full step-wise bid curves for every offer. Each line is represented as the tuple **`{cab_id, period, block_number, step, exclusive_group_index, price, energy, MAV}`**. The parser classifies each row into one of three structures stored under the corresponding unit:
  - `simple_steps` — rows with `block_num == 0`, corresponding to the ordinary simple-offer curve used in the standard EUPHEMIA clearing;
  - `blocks` (simple and exclusive) — rows with `block_num > 0`, grouped by block number, represent block orders; rows with a non-zero exclusive-group index are additionally recorded in `excl_groups` to report exclusive blocks;
  - `SCOs` — the per-period minimum-acceptance volume (MAV) extracted from field 8, which encodes the minimum megawatt dispatch the unit requires in each period if its complex offer is activated. It also includes a fixed term in the CAB file (field "FijoEuro", bytes [54:71]) — a fixed cost in euros that, when greater than zero, independently flags the offer as an SCO regardless of what the MAV field in DET shows.

4.3.3 | Intraday Data Ingestion

The intraday pipeline mirrors the day-ahead pipeline, but operates on the three auction sessions independently. The routines `load_ipibc` and `load_phf` read the marginal-price and matched-program files, respectively, applying the session-start filter so that rows outside the re-auctioned window of a given session are discarded. The routine `parse_icab` reads the intraday offer headers using the official OMIE fixed-width specification, which differs from the day-ahead CAB format in one structurally important way: complex orders in IDAs are restricted to *block orders*, and the minimum-income SCO does not exist. A non-zero value in the `PrcMedBloqOrder` field at positions 59–75 flags the corresponding offer as a block order and stores its required average price. The routine `parse_idet` then reads the intraday detail file with empirical column positions (established by inspection of real files, because the published specification does not always match the delivered layout exactly) and attaches each row either as a floor-price observation or, when the offer number is in the block-order set, as part of a block aggregate.

A critical design decision in this module concerns how incremental intraday activity is measured. For session 1 the baseline against which the matched programme is compared is the day-ahead PDBC; for session 2 it is the session 1 pibc; for session 3 it is the session 2 pibc. Measuring each session against the previous session's final schedule, rather than always against PDBC, isolates the *marginal* volume traded in that specific auction and avoids double-counting energy that was re-confirmed across multiple sessions. An additional filter restricts the computed delta to units that actually appear in IDET, which excludes schedule changes driven by bilateral trades or network corrections rather than by the intraday auction itself. This treatment makes the session-by-session revenue figures directly interpretable as the economic gain attributable to that particular IDA. Figure 4.2 illustrates the resulting baseline structure.

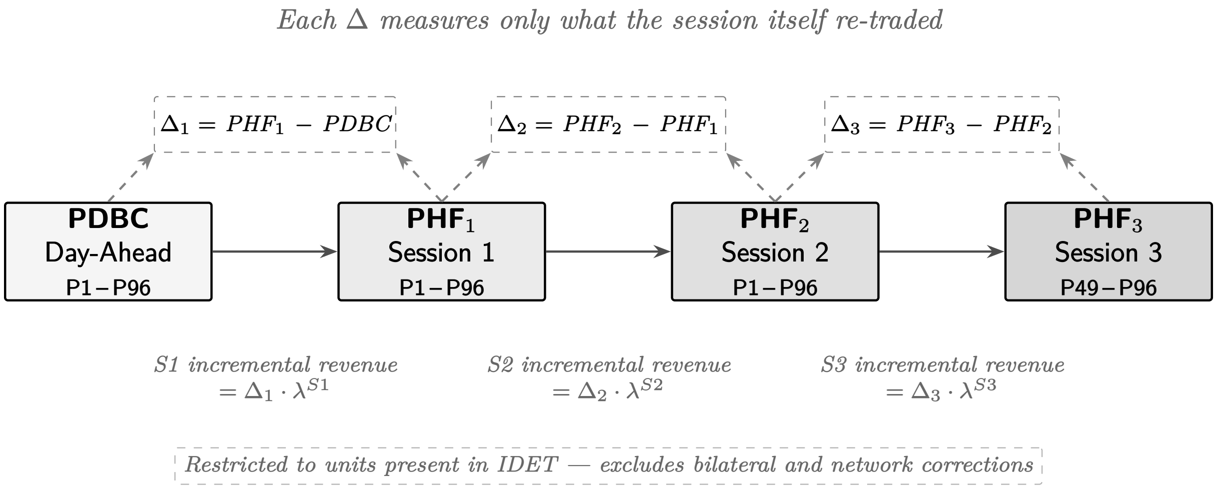

Figure 4.2: Baseline against which each intraday session's incremental volume is measured. Session 1 is compared to the day-ahead schedule (PDBC); sessions 2 and 3 are compared to the preceding session's matched programme (PHF), not back to PDBC. This isolates the marginal volume re-traded in each auction and avoids double-counting energy re-confirmed across sessions. Session 3 covers only the second half of the delivery day (P49–P96).

4.3.4 | Revenue Computation and Verification

Once the files are in memory, the script computes the full set of aggregates needed by the dashboard. Total generation, total revenue, and the generation-weighted average price are computed directly from the nested PDBC dictionary and the price vector. Unit-level net revenue is decomposed into a generation component and a consumption (pumping) component using:

$$\mathrm{net}_i = \sum_{p\in\mathcal{P}} \max\big(0,\, P_{i,p}\big)\, \lambda_p\, \Delta t \;-\; \sum_{p\in\mathcal{P}} \Big|\min\big(0,\, P_{i,p}\big)\Big|\, \lambda_p\, \Delta t, \tag{4.1}$$

where $P_{i,p}$ is the matched megawatt of unit i in period p, λ_p is the Spanish marginal clearing price, and $\Delta t = 0.25\,\text{h}$ under the current quarter-hourly MTU15 design. Expression (4.1) is the formula validated in the backcast section of this thesis, and the dashboard's *Data Integrity Check* tab reports it side by side with a second, independently computed value for the four pumped-hydro pairs (Aguayo, Duero, Sil, La Muela). The two methods agree to within rounding precision on all days tested, which is the core empirical guarantee underlying the optimiser's use of OMIE data.

Complex-order compliance is checked in a dedicated routine: for every unit with a declared minimum income, the script computes the realised revenue, compares it against the threshold, and marks the offer as *activated* (threshold met and unit dispatched) or *withdrawn* (unit did not match). When a MAV is also declared, the routine additionally verifies that the matched megawatt in every MAV-relevant period is at or above the declared floor. Applied to the reference day 1 December 2025, this routine identifies three units with an SCO: the two Ascó nuclear reactors `ASC1` and `ASC2`, both with a MAV condition (994 MW and 867 MW respectively) and a minimum-income threshold met by the realised revenue, and one renewable aggregator `AMRE029` whose minimum-income condition is likewise satisfied. All three SCOs activate cleanly, and all MAV floors hold period by period, which is consistent with the description of EUPHEMIA's treatment of minimum-income conditions given in the state-of-the-art chapter.

4.3.5 | Dashboard Output

The rendering layer is entirely implemented in the string constant `HTML_TEMPLATE`, which contains the full HTML document with an inlined CSS design system, a set of panel containers for the four tabs, and a vanilla-JavaScript controller. The Python side finishes by serializing the aggregated `DATA` dictionary to JSON via `json.dumps`, effectively freezing every numeric quantity into a JavaScript-compatible literal and substituting it into the placeholder inside the template. The JSON payload is the sole bridge between the Python and JavaScript worlds, and it is what makes the HTML file self-sufficient: no network requests to OMIE, no runtime Python interpreter, and no database is needed once the file has been generated.

4.4 | Yearly Block-Order Analyser: `yearly_data.py`

The daily backcast confirms that the parsing pipeline correctly identifies and classifies complex orders across both market sessions, but a single reference day cannot answer the broader question of how block orders are used by the market at scale. To fill that gap, a second tool, `yearly_data.py`, iterates over every ICAB and IDET file across the twelve-month window. It then identifies every submitted block order, and aggregates them by unit, direction, technology, session, and size. The output (`yearly_blocks.html`) is an interactive yearly report whose findings are summarized below.

4.4.1 | Configuration Block

The multi-month script `yearly_data.py` uses a configuration structure suited to a longitudinal study rather than a single-day run. Instead of explicit file paths for each input, the config block exposes four variables. `DATE_FROM` and `DATE_TO` define the analysis window, set to `2025-02-01` and `2026-01-31` in the current run, covering one full year of IDA auction activity. `ID_DIR` points to the root folder containing the intraday files, which the script expects to find organized into monthly subdirectories (`icab_YYYYMM/`, `idet_YYYYMM/`, `phf_YYYYMM/`); `DA_ROOT` points to the folder containing the day-ahead PDBC subdirectories (`pdbc_YYYYMM/`) and defaults to the script's own directory. Two numerical constants complete the block: `N_PERIODS = 96` and `DT = 0.25`, reflecting the quarter-hourly resolution adopted by the Spanish market from October 2025 onwards.

For each calendar day in the date range the script iterates all three IDA sessions, resolving file paths from the date and session index. Re-running the analysis for a different period therefore requires changing only the two date strings.

4.4.2 | Identification Rule

A block order is identified by a single, unambiguous rule taken from the official OMIE specification: an offer at positions 13–20 of an ICAB line is a block order if and only if the field `PrcMedBloqOrder` at positions 59–75 (an `F17.3` real) is non-zero (which means that `PorcMinBloqOrder` at positions 76–80 is also non-zero). No heuristic based on the `TipoNec` field or the offer name is used; the rule is purely syntactic. The value of `PrcMedBloqOrder` is the required average price that the block must earn across the periods it covers in order to be accepted by EUPHEMIA, and the companion field `PorcMinBloqOrder` at positions 76–80 encodes the minimum fill percentage. The accompanying IDET file is then inspected to recover the periods covered by each block and the requested megawatt, which is averaged across the block's periods to produce a single representative size. Sizes are bucketed into three classes: small (< 50 MW), medium (50–200 MW), and large (> 200 MW).

Unlike in the day-ahead, here generators can buy or sell. The direction of the orders thus becomes a new important element to analyse. Direction is read from `cv` at position 54 of the same ICAB line: `V` identifies a *venta* (a sell offer, submitted by a generator), and `C` identifies a *compra* (a buy offer, submitted either by a consumer, a pump, or a trader). Both directions are retained, because the complex-offer product is symmetric and pumped hydro in particular uses the buy side extensively.

4.4.3 | Aggregation Pipeline

The main loop walks every calendar day of the configured year and, for each day, every session from 1 upwards. For each file pair it calls `parse_icab_blocks` (which returns the set of block orders together with the unit metadata) and `parse_idet_blocks` (which returns the periods and megawatt profile for each block). Every successfully parsed block order becomes one record in a flat list, and the aggregation step builds five derived objects: a per-unit summary, a technology breakdown, a session breakdown, a direction–size cross-tabulation, and a timeline of submissions per day. A heatmap of the top 20 units against the calendar of the month is constructed separately so that the intensity of block-order usage can be inspected at a glance. All derived objects are serialized to JSON and dropped into the HTML template in the same way as the daily backcast.

The analyser is designed to answer two prior questions before the battery model is specified. The first is empirical: whether intraday block orders are an active product in the Spanish market or one used intermittently. The second is structural: whether the units that use them map onto assets with multi-period operating constraints (such as start-up costs and minimum-run windows for thermal plant or round-trip efficiency and reservoir coupling for pumped hydro) or whether usage patterns bear no analogy to a dispatchable storage asset.

A battery system has state-of-charge dynamics, round-trip efficiency, and charge/discharge coupling across periods that are structurally closer to (both pumped and storage) hydro than to either a combined cycle or a thermal unit. If intraday block orders are already the instrument of choice for assets with precisely this kind of multi-period coupling, the question is whether the existing order shape is expressive enough to capture a battery's operating constraints or whether the value accessible in IDAs would be higher under an extended order type that explicitly represents a charge-and-discharge path. The empirical findings that answer the first two questions are reported in Section 5; the battery model that follows from those findings is developed in Section 4.6.

4.5 | Price Forecasting Model: LEAR-LASSO Framework

Participation in OMIE's intraday auctions requires submitting orders before market closure, when the session's clearing prices are not yet known. The battery must therefore bid on a forecast of the clearing price rather than the realised price. A LEAR (Lasso Estimated AutoRegressive) model is implemented for this purpose, providing the exogenous price signal $\hat{p}_{d,t}$ used in the order-type optimisation models defined in Section 4.8.

LEAR was chosen over gradient boosting and neural network alternatives because it matches their point-forecast accuracy on OMIE-scale datasets [33] with far less setup. The only tuning decision is how strongly to penalise unhelpful predictors, which cross-validation selects automatically. The training window starts at 35 days and expands to roughly 430 days; the built-in penalty drops irrelevant features on its own, so the model performs well even on short samples where neural networks would start fitting noise. The per-period coefficients are also interpretable, an advantage when the methodology must be explained alongside the results.

LEAR fits one LASSO-regularised regression per quarter-hourly period rather than a single model across all 96 slots. Each estimator trains only on the prices of its own slot, so the 08:00–08:15 regression learns morning-peak dynamics without being contaminated by overnight prices [35], [43]. Pre-October 2025 files have 24 hourly prices; these are repeated four times per hour to fill the 96-slot grid, so every estimator sees a consistent quarter-hourly input throughout the full sample.

The 18 input features follow the standard LEAR specification of [33] with the addition of a monthly seasonality pair: five same-period lagged prices ($d-1$, $d-2$, $d-3$, $d-7$, $d-14$), three daily mean prices, yesterday's price range (min and max), the seven-day same-period rolling mean, and six cyclical calendar encodings (day-of-week, month, and intraday position).

Same-slot lags (Six features). Five lagged prices from the same time slot on previous days: yesterday ($p_{d-1,t}$), two days ago ($p_{d-2,t}$), three days ago ($p_{d-3,t}$), last week ($p_{d-7,t}$), and two weeks ago ($p_{d-14,t}$). Same-slot lags outperform adjacent-slot lags because the 08:00 price is better predicted by last Monday's 08:00 than by yesterday's 07:45. The three consecutive-day lags track short-run price momentum. The weekly lag captures the weekday–weekend pattern. The two-week lag adds a second pass at the weekly cycle and improves out-of-sample accuracy in standard EPF benchmarks [43]. The sixth feature is the 7-day average of the same slot, $\frac{1}{7}\sum_{k=1}^{7} p_{d-k,t}$, which tracks the recent price level for that period and dampens the effect of one-off spikes.

Daily-level indicators (Five features). Three daily average prices ($\bar{p}_{d-1}$, $\bar{p}_{d-2}$, $\bar{p}_{d-7}$) capture market-wide shifts: a day when gas is expensive or hydro is scarce tends to lift all 96 periods together, and same-slot lags alone miss that. Yesterday's within-day minimum and maximum complete the group. A wide spread the day before signals a volatile market where order-type choice matters; a narrow spread signals low arbitrage potential.

Calendar encodings (Six features). Sine–cosine pairs encode three cycles: day of week ($2\pi/7$), month of year ($2\pi/12$), and period within the day ($2\pi/96$). Paired trig functions rather than dummy variables mean Sunday and Monday sit next to each other in the feature space, as do December and January — the calendar wraps without a discontinuity at the boundary.

LASSO adds a penalty that zeros out unhelpful coefficients, so the model automatically drops features that do not improve the forecast. The penalty strength α is chosen by testing different values on past data and picking the one that performs best out-of-sample. Forecast accuracy is reported in the results chapter as MAE: the mean absolute error in €/MWh, averaged across all periods and all evaluation days in the 14-day walk-forward window.

Day-ahead price anchor (One feature). The eighteenth and final feature is the day-ahead clearing price for the same calendar date and period, $p^{DA}d, t$. Unlike the other seventeen, which describe historical price behaviour, this feature is a same-day, forward-looking value: the day-ahead auction for date d clears before any of date d's three intraday sessions open, so $p^{DA}d, t$ is genuinely known in advance and its inclusion introduces no look-ahead bias. It functions as a regime anchor that has a same-day reference point that the same-slot lags and daily-level indicators, all drawn from prior days, cannot provide. It grounds the forecast in the market's own most recent, most relevant price signal for that exact period. This feature is specific to the intraday forecaster; the day-ahead LEAR model, which had no earlier same-day auction to draw on, used only the seventeen features described above.

Walk-forward validation. Model performance is evaluated using a walk-forward (rolling-origin) scheme following the approach of [33]. The training window is *expanding*: the first forecast is produced on day 35 of the sample, determined by the minimum training requirement of 21 days beyond the longest autoregressive lag of 14 days ($21 + 14 = 35$), after which the window grows by one day with each step and no historical observations are ever discarded. A fresh set of 96 period-specific LASSO models is re-estimated daily, so the forecast for each target day is produced by a model that has never seen that day's prices. The IDA data covers Feb 2025 – May 2026; the first 35 days of each session's sample are held out as burn-in, since this window combines the 14-day longest lag with a 21-day minimum training buffer, the latter ensuring at least three cross-validation folds are available for LassoCV alpha selection. The expanding design is preferred over a fixed rolling window because electricity price series carry long-run seasonality signals that remain informative well beyond a single year, and LASSO's ℓ_1 regularisation controls overfitting as the dataset grows, rendering the accumulation of older observations cost-effective rather than detrimental.

A separate LEAR model is estimated for each of the three intraday auction sessions, using that session's realised `marginalpibc` price series as the target variable. The 96-period forecast for session s is produced independently of the other sessions, with features constructed exclusively from prices observed prior to session s's market close. This means the model for session 2 can in principle use session 1's realised prices as an additional feature, but the present implementation uses only day-level lags for simplicity and treats the three session forecasts as parallel rather than sequential.

4.6 | BESS Welfare Benchmark Model

The welfare benchmark model computes W^*: the maximum net arbitrage revenue achievable by a perfectly informed, price-taking battery operating in the OMIE intraday market, after deducting the same degradation (which will be discussed later) cost c_{deg} applied in the order-type LPs. It serves as an upper bound against which all bidding strategies are evaluated. The gap $W^* - W_{\text{strategy}}$ therefore isolates the pure structural cost imposed by the order type's bidding constraints: since both sides are net of the same degradation rate, degradation cancels out and the gap reflects only the revenue lost because the submitted order cannot replicate the perfect-foresight LP dispatch.

4.6.1 | Battery Assumptions

Table 4.2: Battery system assumptions

Parameter	Symbol	Value	Rationale
Energy capacity	E_{cap}	10 MWh	Mid-size utility-scale BESS
Power rating	$P_{\max}$	5 MW	2-hour duration (standard arbitrage setup)
C-rate	—	0.5 C	Derived ($P_{\max}/E_{\text{cap}} = 5/10$)
Degradation cost	c_{deg}	10 €/MWh	Per-MWh cost charged on each unit discharged
Charge/discharge eff.	η_c, η_d	0.92 each	Industry-standard per half-cycle efficiency
Round-trip eff.	η_{RT}	84.6 %	$\eta_c \times \eta_d = 0.92^2$
SoC range	$S^{\min}$–$S^{\max}$	5 %–95 %	Lifetime-preserving bounds
Initial SoC	$s[0]$	$S^{\min} E_{\text{cap}} = 0.05 \times 10\,\text{MWh} = 0.5\,\text{MWh}$	Both strategies start at minimum SoC
Throughput cap	$E_{\max}$	18 MWh	2-cycle regime

The 2-hour discharge duration corresponds to a C-rate of 0.5 C and is consistent with the shortest standard configuration covered by NREL's Annual Technology Baseline for utility-scale battery storage [44]. This duration is appropriate for intraday energy arbitrage: a slower 0.25 C configuration (four-hour duration) would reduce the asset's ability to capture price spreads within OMIE's inter-session windows, while discharge rates above 1 C are associated with short-duration ancillary services such as primary frequency response and regulation, which operate on timescales of seconds to minutes rather than hours [45]. The 84.6 % round-trip efficiency and 90 % usable capacity together reflect a conservative but commercially realistic LFP configuration consistent with Spanish merchant storage projects currently in development.

The degradation cost is modelled as a flat per-MWh charge on discharged energy. Lithium-ion cell degradation is physically a non-linear process dependent on cycle depth and charge/discharge history [46], but tracking this directly (e.g. via cycle-counting algorithms) is not compatible with a linear programme. Following the approach used in comparable battery-scheduling optimisation [47], this thesis instead normalises expected battery lifetime into total energy throughput and prorates the cell replacement cost into a flat per-MWh coefficient applied to charged and discharged energy; this is an explicit simplification of the underlying degradation physics, adopted specifically to preserve LP tractability.

It differs from the cycle-counting degradation model used in a comparable European market study by Nuñez et al. [48], who track discrete charge/discharge cycles and model capacity fade of up to 20 % after 5,000 cycles; that approach requires a mixed-integer formulation, whereas the linear per-MWh cost preserves tractability across the full Monte Carlo scenario evaluation used here (300 scenarios × 96 periods per session, evaluated for every order type on every day of the sample). The choice comes at the cost of not penalising dispatch fragmentation independently of total energy throughput: five 2 MWh charges cost exactly the same in degradation as one 10 MWh charge of equal total energy. The dispatch charts in Chapter 6 should be read with this in mind. A model incorporating cycle-count-based degradation would likely favour coarser, less fragmented dispatch and report somewhat lower Simple-strategy revenue than found here.

Profitability condition. The round-trip efficiency $\eta_{\mathrm{RT}} = 0.92^2 \approx 0.846$ sets a minimum price spread before a charge–discharge cycle is profitable. A battery that buys 1 MWh of energy recovers only η_{RT} MWh of it after conversion losses, so the discharge price must clear not the charge price itself but the charge price inflated by $1/\eta_{\mathrm{RT}}$, plus the degradation cost charged per MWh discharged. The full break-even condition is:

$$p_{\mathrm{sell}} > \frac{p_{\mathrm{buy}}}{\eta_{\mathrm{RT}}} + c_{\mathrm{deg}}$$

At a representative charge price of 40€/MWh and $c_{\mathrm{deg}} = 10$€/MWh, the required discharge price is $40/0.846 + 10 \approx 57.3$ €/MWh — a required spread of 17.3,€/MWh, equivalent to 43.1% of the charge price itself. Whether this threshold is met in practice will be evaluated in the results section for each session.

4.6.2 | Optimisation Problem

The welfare benchmark is formulated as a Linear Programme (LP) over 96 quarter-hourly periods, using the realised intraday clearing prices as the price signal. Throughout this model, periods are indexed $t = 1, \dots, 96$, with $t = 1$ corresponding to the interval 00:00–00:15 CET and $t = 96$ to 23:45–00:00 CET. This LP differs from Simple's own dispatch LP (Section 4.6) in exactly one respect: the price signal it is solved against. Simple's LP is solved on the LEAR forecast $\hat{p}[t]$ before the session opens, and its resulting fixed schedule is only settled against realised prices afterward; the welfare benchmark is solved directly on the realised price $p[t]$, representing a battery with perfect foresight; so it is the same LP formulation, but applied to a different, ex-post price vector.

Objective function:

$$\max \sum_{t} \Big[\Big(p[t] - c_{\text{deg}}\Big)\, x_d[t]\, \Delta t - p[t]\, x_c[t]\, \Delta t \Big] \tag{4.2}$$

Before the constraint set, it is worth fixing precisely what each symbol represents:

- $p[t]$ — the realised intraday clearing price for period t (the actual marginal value, known only after the session clears).
- $\Delta t = 0.25$ h — the settlement period length (15 minutes), converting the power variables into energy wherever multiplied by Δt.
- c_{deg} — the degradation cost, 10 €/MWh discharged, charged only on energy actually discharged, not on energy charged.

Decision variables:

Table 4.3: Decision variables of the welfare benchmark LP

Variable	Type	Meaning
$x_c[t]$	continuous	Charging power (MW)
$x_d[t]$	continuous	Discharging power (MW)
$s[t]$	continuous	State of charge (MWh)

$x_c[t]$ and $x_d[t]$ are the battery's charging and discharging power in period t (MW); by construction only one is non-zero in an optimal solution, enforced by the non-simultaneity constraint below. $s[t]$ is the state of charge at the end of period t (MWh), linked across periods by the SoC-dynamics equation.

State-of-charge dynamics:

$$s[t] = s[t-1] + \eta_c\, x_c[t]\, \Delta t - \frac{x_d[t]\, \Delta t}{\eta_d} \tag{4.3}$$

with $s[0] = S^{\min} E_{\text{cap}}$ for the welfare benchmark (starts at minimum state of charge).

where η_c, η_d are the charging and discharging efficiencies (0.92 each), together giving the round-trip efficiency $\eta_{\text{RT}} = \eta_c \eta_d \approx 0.846$; $S^{\min} E_{\text{cap}} = 0.5$ MWh (5% of the 10 MWh capacity) is the SoC floor the battery both starts and ends every session at, so the battery opens and closes each session empty rather than carrying energy value across session boundaries.

Constraints:

- *SoC bounds:*

$$S^{\min} E_{\text{cap}} \leq s[t] \leq S^{\max} E_{\text{cap}}$$

  where $S^{\min} = 5\%$, $S^{\max} = 95\%$, $E_{\text{cap}} = 10$ MWh. These bounds keep the battery away from full and empty extremes, where degradation per cycle accelerates sharply — the depth-of-discharge is one of the main drivers of capacity fade in lithium-ion cells, alongside state of charge and cycling frequency.

- *Power limits:*

$$0 \leq x_c[t] \leq P_{\max}, \qquad 0 \leq x_d[t] \leq P_{\max}$$

  where $P_{\max} = 5$ MW is the battery's rated power, the physical ceiling the inverter and cell chemistry can sustain per period.

- *Non-simultaneity:*

$$x_c[t] + x_d[t] \leq P_{\max}$$

  Written as a continuous inequality rather than a hard exclusivity constraint on binary variables, this is the LP relaxation of mutual exclusion. It is not merely a convenient approximation: under price-taking behaviour, a rational optimiser never benefits from charging and discharging within the same period at the same clearing price, since doing so only destroys energy through round-trip losses for no revenue gain, so the relaxed optimum coincides exactly with what a binary exclusivity constraint would enforce. This is what allows the welfare benchmark, like Simple, to remain a fast LP rather than requiring the binary variables the Loop Order's paired-block structure does (Section 4.8.2).

- *Daily throughput cap:*

$$\sum_t x_d[t]\, \Delta t \leq E_{\max}$$

  where $E_{\max}$ is 18 MWh (2-cycle regime, the value used throughout this backtest). Ampere-throughput is itself a recognised degradation driver independent of depth-of-discharge, so this cap keeps the benchmark's daily cycling within the same physical envelope assumed for the other strategies, rather than letting perfect foresight extract value through unlimited cycling.

- *Terminal neutrality:*

$$s[96] = S^{\min} E_{\text{cap}}$$

  Forcing the battery to close each session at the same SoC floor it opened at prevents the LP from ending the day in an advantageous state of charge that would not be available at the start of the next session, which would otherwise inflate the benchmark's profit artificially.

4.6.3 | Model Scope and Simplifications

Included in scope.

- Energy arbitrage revenue
- Efficiency losses
- Intertemporal coupling (SoC dynamics)

Excluded by design.

- Auxiliary consumption
- Grid constraints
- Ancillary services revenue

These exclusions are deliberate. The model is a welfare benchmark rather than a financial forecast: its purpose is to isolate the maximum value available from energy arbitrage from the day-ahead electricity market alone, so that any shortfall in the bidding strategies evaluated later can be attributed unambiguously to order-type constraints rather than to omitted cost terms.

4.6.4 | Solution Process

For each day in the sample the model loads the 96-period OMIE intraday price series, instantiates the battery parameters, builds the LP, and solves it with the HiGHS solver (via `scipy.optimize.linprog`), consistent with the Simple Order LP described in Section 4.6.2. The optimal dispatch schedule (x_c^*, x_d^*) and the resulting objective value W^* are extracted and stored. Typical solve time is under one second per day.

4.6.5 | Interpretation of W^*

W^* is the revenue that a battery would earn if it had perfect foresight of intraday prices, faced no bidding constraints, and behaved as a price-taker. No realistically implementable strategy can exceed it. The welfare gap $W^* - W_{\text{strategy}}$ therefore measures the structural cost imposed by each order type: a small gap indicates that the order type is expressive enough to capture most of the available arbitrage value, while a persistently large gap indicates that the market design itself restricts economically efficient battery participation and motivates the introduction of a richer order type.

Throughout, W_{strategy} denotes the *realised* revenue earned by executing the order type chosen ex ante, based on forecast prices and the 300-scenario expected revenue criterion, evaluated ex post against the actual IDA clearing prices from `marginalpibc`. The gap therefore reflects the revenue lost relative to perfect foresight, attributable to forecast error, order-type constraints, and degradation costs as discussed above.

4.6.6 | Gap Aggregation and Reporting Methodology

The welfare gap $W^* - W_{\text{strategy}}$ is computed for each day and session in the sample. Chapter 5 (results) reports this gap primarily through its mean (aggregated per session and combined across the 387-day sample) and its relative form, $(W^* - W_{\text{strategy}})/W^*$, expressed as a percentage of the oracle benchmark; Section5.3.7 additionally decomposes the gap by root cause (forecast failure type, loss-day contribution). All figures reported throughout Chapter 5 use the two-cycle throughput regime ($E_{\max} = 18$ MWh), the default regime this backtest was run under.

4.7 | Monte Carlo Scenario Generation

The LEAR point forecast described in Section 4.5 provides the central price trajectory $\hat{p}_{d,t}$ for each period $t \in \{1, \dots, 96\}$ of target day d. To capture price uncertainty before market close, the model augments the point forecast with a set of $S = 300$ independent price scenarios. $S = 300$ was chosen as a conservative threshold well above the $1/\alpha = 10$ minimum for stable VaR estimation, while remaining within a feasible single-machine runtime.

4.7.1 | Residual estimation

Out-of-sample forecast errors are collected on a rolling window of up to 14 days immediately preceding the target date. For each period t, given the residual $r_{d-k,t} = p_{d-k,t} - \hat{p}_{d-k,t}$ over the N most recent evaluable days ($N \leq 14$), the residual standard deviation is estimated as:

$$\hat{\sigma}_t = \sqrt{\frac{1}{N}\sum_{k=1}^{N}\left(r_{d-k,t} - \bar{r}_t\right)^2}, \qquad \bar{r}_t = \frac{1}{N}\sum_{k=1}^{N} r_{d-k,t} \tag{4.4}$$

$\hat{\sigma}_t$ is clipped to a floor of 1 €/MWh. The window N is capped at 14 but shrinks near the start of the sample, when fewer prior days are available, down to a minimum of 3. Below that, there isn't enough data to estimate volatility per period, so $\hat{\sigma}_t$ falls back to a flat 5 €/MWh. This fallback only applies during the first few days of the sample, before the burn-in period ends.

The choice of window length is a trade-off. A shorter window adapts faster to structural breaks like the October 2025 shift to quarter-hourly resolution, but is noisier, since it's more exposed to atypical days. A longer window smooths out that noise but risks carrying stale price dynamics forward.

4.7.2 | Scenario draws

Conditional on $\hat{p}_{d,t}$ and $\hat{\sigma}_t$, each scenario price is drawn independently across periods from a Gaussian truncated at zero:

$$p_{d,t}^{(s)} \sim \mathcal{N}\left(\hat{p}_{d,t},\, \hat{\sigma}_t^2\right)\Big|_{p \geq 0}, \quad s = 1, \dots, 300;\ t = 1, \dots, 96 \tag{4.5}$$

The model treats periods as independent (no copula structure). This simplification ignores temporal error correlation but ranks order types correctly by expected revenue before market close, consistent with the LEAR literature [35]. The resulting scenario matrix $\mathbf{P} \in \mathbb{R}^{300 \times 96}$ feeds all order-type evaluation routines described in Section 4.8.

The Gaussian assumption understates tail probabilities relative to the empirical distribution of OMIE intraday prices, which exhibit positive skewness and excess kurtosis during renewable surplus events. A separate and more consequential limitation is the independence assumption itself. Scenarios are drawn independent and identically distributed per period with no cross-period correlation structure: the price realisation at 18:00 has no bearing on the draw at 18:15. In reality, intraday price dynamics are strongly autocorrelated.

For example, a price spike at 18:00 makes an elevated price at 18:15 substantially more likely, and the sustained multi-hour ramps that characterise high-value arbitrage days are a direct expression of this temporal dependence. Independence suppresses exactly these joint realisations: under i.i.d. draws, a scenario where prices remain elevated for six consecutive periods is far rarer than it is in practice, because each period is drawn without reference to its neighbours.

This limitation was tested directly rather than left as a theoretical concern. An alternative scenario generator was implemented, drawing each scenario by resampling a whole historical residual-day vector (from the same rolling 14-day window used to estimate $\hat{\sigma}_t$) rather than synthesising independent per-period noise. This preserves whatever real cross-period correlation and fat-tailed behaviour the historical errors actually exhibit, at the cost of drawing from only the N error patterns observed in the recent window rather than a continuous distribution. Re-running the full backtest with this bootstrap-based scenario matrix in place of the Gaussian one leaves every headline result in this thesis unchanged: neither Simple's nor rule-based Loop's realised revenue, the figures reported as this thesis's primary results throughout Chapter 6, depend on the scenario matrix at all, so neither can be affected by any choice of scenario-generation method. The Gaussian draw in Equation (4.5) is retained throughout this thesis for its simplicity, since the alternative was confirmed not to change any result this thesis reports.

4.8 | Order-Type Optimisation Models

OMIE's intraday auction sessions accept two order types currently available to market participants: Simple Orders and Simple Block Orders. This model evaluates Simple Orders and the Loop Order (C88), a paired charge/discharge block structure available in EPEX SPOT sessions and proposed here as an addition to OMIE's IDA product set, together with a NoBid option, which represents rational abstention when all evaluated strategies yield negative expected revenue across scenarios. All order types are evaluated against IDA clearing prices from `marginalpibc`, using the session structure described in Section 4.4.

Block Orders, available in OMIE's IDA sessions and actively used by thermal and pumped-hydro units as documented in Section 5.2, are excluded from the order-type comparison a priori rather than evaluated head-to-head against Simple and Loop. A Block Order commits the battery to a fixed power in a single direction with no mechanism to pair charge and discharge legs, making it structurally weaker than both the LP-optimised Simple schedule (which can reposition every 15 minutes) and the Loop Order (which uses blocks, but actually pairs the charge and discharge legs) for a cycling asset such as a BESS. This structural dominance argument, rather than a numerical backtest comparison, is the basis for restricting the formal selection criterion in equation (4.13) to Simple, Loop, and NoBid.

Session independence assumption. The three intraday sessions are treated as independent optimisation problems. At the start of each session the battery's state of charge is reset to the minimum charge level ($s[0] = S^{\min} E_{\text{cap}}$, i.e. 5 % of capacity), avoiding cross-session SoC dynamics that the model does not track. This simplification allows the order-type comparison to be conducted cleanly within a single session. Extending the framework to a sequential model, where the terminal SoC of session s becomes the initial condition of session $s+1$ and the bidding decision in each session is conditioned on the expected residual value of remaining sessions, is a natural direction for future work. Within the scope of this thesis, the contribution is the order-type comparison itself, not the cross-session scheduling problem.

4.8.1 | Simple Order: Linear Programme

The Simple Order dispatches the battery period-by-period through an LP solved on the LEAR point forecast. Given orders clear at the marginal price, the problem is deterministic given $\hat{p}_{d,t}$, and no binary variables are required.

Decision variables.

- $x_c[t] \geq 0$: charging power (MW), $t = 1, \ldots, 96$.
- $x_d[t] \geq 0$: discharging power (MW).
- $s[t] \geq 0$: state of charge (MWh).

Objective. Maximise arbitrage revenue net of degradation:

$$\max \sum_{t=1}^{96} \Big[(\hat{p}_{d,t} - c_{\text{deg}})\, x_d[t] - \hat{p}_{d,t}\, x_c[t] \Big]\, \Delta t \tag{4.6}$$

where $\Delta t = 0.25\,\text{h}$ (MTU15 quarter-hourly convention) and $c_{\text{deg}} = 10\,€/\text{MWh}$ is the per-MWh degradation cost charged on each unit of energy discharged.

Constraints. State-of-charge dynamics follow equation (4.3); the initial SoC for each session is $s[0] = S^{\min} E_{\text{cap}}$ (5 %), as established in the session independence assumption at the opening of Section 4.8.

SoC operating bounds:

$$S^{\min} E_{\text{cap}} \leq s[t] \leq S^{\max} E_{\text{cap}}, \qquad S^{\min} = 0.05, \quad S^{\max} = 0.95 \tag{4.7}$$

Power limits:

$$0 \leq x_c[t] \leq P^{\max}, \quad 0 \leq x_d[t] \leq P^{\max}, \qquad P^{\max} = 5\,\text{MW} \tag{4.8}$$

Non-simultaneity (LP relaxation):

$$x_c[t] + x_d[t] \leq P^{\max} \tag{4.9}$$

Under perfect price foresight and price-taking behaviour, the LP relaxation of the mutual-exclusion constraint yields an identical optimal solution to the binary formulation ($u_c[t] + u_d[t] \leq 1$, replaced here by $x_c[t] + x_d[t] \leq P_{\max}$), since a rational optimiser will never simultaneously charge and discharge at the same period price. The Loop Order model, by contrast, requires binary variables to enforce the paired block commitment structure; see Section 4.8.2.

Daily throughput cap:

$$\sum_{t=1}^{96} x_d[t]\,\Delta t \le E_{\max} \tag{4.10}$$

where $E_{\max}$ equals 9 MWh (one-cycle regime) or 18 MWh (two-cycle regime). Terminal neutrality:

$$s[96] = S^{\min}\,E_{\text{cap}} \tag{4.11}$$

No ramp-rate or minimum on-time constraint is included in the LP. Grid-scale LFP batteries transition from zero to rated power in well under one second [45], which is non-binding relative to the 15-minute settlement period. No equivalent regulatory requirement has been identified for BESS participating in OMIE's wholesale intraday market, and comparable studies on the Iberian market do not model such a constraint [48], [49]. HiGHS (via `scipy.optimize.linprog`) solves the LP. The resulting schedule (x_c^*, x_d^*) is then evaluated across all 300 price scenarios to build $\mathbb{E}[R^{\text{Simple}}]$ and its empirical distribution.

4.8.2 | Loop Order (C88)

The Loop Order (C88) pairs a charge block over $[t_c^s, t_c^e]$ with a discharge block over $[t_d^s, t_d^e]$, subject to no temporal overlap $(t_d^s > t_c^e)$. Currently available in EPEX SPOT sessions, it is evaluated here as a proposed addition to OMIE's IDA product set. The analysis should therefore be interpreted as a policy counterfactual: what revenue would a C88-equivalent instrument generate if introduced in OMIE under these clearing rules, not as a prediction of actual OMIE outcomes. The clearing condition, joint acceptance only if the realised average sell price exceeds the bid sell price and the realised average buy price falls below the bid buy price, is a simplified representation of block matching, abstracting from the aggregate order-book effects that EUPHEMIA's optimiser would account for.

The model chooses the Loop Order only if the expected net cycle revenue for the battery across all $S = 300$ scenarios is positive. The per-scenario net revenue is:

$$R^{\text{net},(s)} = P^{\max}\,\Delta t \sum_{t\in\mathcal{D}} p_{d,t}^{(s)} - P^{\max}\,\Delta t \sum_{t\in\mathcal{C}} p_{d,t}^{(s)} - c_{\text{deg}}\,E_d \tag{4.12}$$

where $\mathcal{C} = \{t : t_c^s \le t \le t_c^e\}$ and $\mathcal{D} = \{t : t_d^s \le t \le t_d^e\}$ are the charge and discharge index sets, $E_d = P^{\max} \cdot (t_d^e - t_d^s + 1) \cdot \Delta t$ is the discharge energy, and $p_{d,t}^{(s)}$ is drawn from the scenario matrix defined in (4.5). The battery submits when $\mathbb{E}_s[R^{\text{net},(s)}] > 0$; individual scenario revenues may be negative. Once submitted, actual clearing depends on whether this condition is met by the realised prices. When it fails, the order is rejected and the battery earns nothing.

C88's clearing condition is a declared-price-limit test, not a break-even test: a leg clears once realised prices beat the operator's own declared bid, independent of whether the resulting margin actually covers the battery's degradation cost, so a Loop pairing built on an incorrect forecast can clear exactly as designed and still realise a loss. This is not the only clearing philosophy available in this market. OMIE's own day-ahead rules already define a structurally different mechanism for a comparable purpose: the Minimum Income Condition (MIC, Article 28.1.2.1 of the day-ahead and intraday market rules), under which a seller may attach to a production unit's bid the condition that the bid is only considered submitted if it earns at least a declared minimum income, a fixed term in Euros plus a variable term per MWh, evaluated against the unit's realised income for the day, with the bid rejected outright, rather than settled at a loss, if the condition is not met. MIC is applied today to conventional generation, not to paired storage legs, but its existence shows that a reject-on-loss clearing rule is not a theoretical alternative to C88's price-limit test but an established category already cleared daily in this market. A natural, if secondary, refinement of the Loop instrument studied here would therefore be a clearing condition modelled on MIC rather than on C88's declared-price limits: instead of testing realised prices against the operator's own bid, each leg (or family of tied legs) would clear only if its realised margin, evaluated at actual prices, covers the battery's degradation cost. Such a condition would not change which windows are searched or selected; the scheduling logic studied throughout this chapter is unaffected. It would only change whether a given schedule is allowed to settle once its outcome is known. Section 5.3.12 bounds this refinement's effect using data already in this backtest, without implementing the mechanism itself.

The candidate generation iterates over a fixed menu of block lengths (2, 4, 8, 12, 16, and 24 quarter-hourly periods, corresponding to 0.5–6 h), but every candidate is filtered by an energy-feasibility constraint before scoring: the product of rated power, block length, and period duration must not exceed the per-cycle energy cap ($P^{\max} \times L \times \Delta t \leq E_{\text{cycle}}$). At the current parameterisation ($P^{\max} = 5\,\text{MW}$, $\Delta t = 0.25\,\text{h}$, $E_{\text{cycle}} = 9\,\text{MWh}$), the longest block that survives this filter at full rated power is $9/(5 \times 0.25) = 7.2$ periods, so in practice only lengths of 2 and 4 periods (0.5 h and 1 h) are energy-feasible; longer nominal lengths are rejected outright rather than derated to a lower power level. The menu therefore defines a search space from which the energy cap selects the physically deliverable subset, not a set of durations the battery actually operates at full power for.

In practice, the shortlisted pairs consistently align with the two highest-spread intervals in the OMIE intraday data: overnight charge paired with morning discharge, and midday solar-driven charge paired with evening discharge.

4.8.3 | Rule-Based Loop Scheduling: An Alternative to the Grid Search

The grid search above ranks candidates by their scenario-based expected revenue, which structurally favours the shortest block whose average price is most extreme rather than the block that trades the most energy at a still-profitable price: because a longer block necessarily incorporates periods further from the single cheapest or most expensive instant, extending a block's length pulls its average price toward the session mean and lowers its ranking score even when the additional volume remains profitable. Chapter 6 shows this causes the grid search to systematically under-use the battery's own capacity, typically settling on a 4-period charge paired with a 2-period discharge regardless of how much cheaper or more expensive capacity remains available in the same trough or peak.

A second, independent Loop scheduling method was implemented to test whether a computationally lighter approach could recover more of this foregone capacity. Rather than exhaustively scoring every block-length combination against 300 Monte Carlo scenarios, the rule-based method works directly from the LEAR point forecast $\hat{p}_{d,t}$: it searches for the cheapest contiguous window of up to a fixed maximum duration of 16 quarter-hourly periods (four hours), further clamped by the per-cycle energy cap to at most 7 periods (1.75 h) under this battery's default parameters, so the nominal four-hour ceiling is not reached in practice. It pairs this window with the highest-value contiguous discharge window occurring strictly after it, and trims whichever side is oversized relative to the other so that declared charge and discharge energies are mutually consistent under the round-trip efficiency constraint ($E_d \leq \eta_c\eta_d\, E_c$) by construction, rather than by post-hoc capping. The candidate pair is scored directly by forecast net profit (equation (4.12) evaluated on $\hat{p}_{d,t}$ rather than on the scenario matrix) and submitted only if this forecast profit is positive; settlement follows the same C88 net-margin condition as the grid-search method, evaluated against realised prices.

Unlike the grid search, which searches for a second cycle only in the span strictly after the first cycle's discharge window ends, the rule-based method searches *both* remaining spans: strictly before the first cycle's charge window begins, and strictly after its discharge window ends, keeping whichever candidate scores higher on forecast profit. Restricting the search to only the after-discharge span, as an earlier version of this method did, silently discards a genuine earlier-in-day opportunity whenever it does not conflict with the first cycle's state-of-charge requirements; since each candidate span still requires the leg to be fully contained within it, real-time state-of-charge feasibility is preserved regardless of which cycle occurs first chronologically. The grid-search method, by contrast, only ever searches after the first cycle's discharge window (Chapter 6), which is one further reason its Loop revenue trails the rule-based method's.

This design is motivated by two results from the wider battery-scheduling literature rather than developed from first principles. Becchi et al. (2024) show that a simple, threshold-based rule-scheduling algorithm, evaluated with no forecast optimisation and no Monte Carlo step, closely approaches LP-optimal battery revenue across a range of test scenarios, with the two methods' revenue curves reported as visually indistinguishable at the scale of their comparison plots; the efficiency gain comes from avoiding an expensive search, not from a more sophisticated objective [50]. Modo Energy's (2023) industry analysis of GB continuous intraday battery arbitrage separately finds that full-duration, two-cycle-per-day operation is the revenue-maximising physical configuration for battery participation in short-term power markets [51]. Neither result specifies a scheduling rule for a paired block-clearing instrument like Loop C88 directly, but both motivate testing whether a method that explicitly targets full-duration capacity utilisation, rather than the sharpest instantaneous spread, performs better for this specific instrument than the exhaustive scenario-ranked search. Chapter 6 reports both methods' backtest results side by side and finds the rule-based method achieves substantially higher realised Loop revenue in every session; the results presented as the primary Loop backtest throughout Chapter 6 use the rule-based method, with the grid search retained as a documented point of comparison.

4.8.4 | Order selection criterion

Each day the headline order-type comparison selects the order type with the highest *realised* revenue among Simple, Loop, and NoBid, evaluated on actual settlement prices rather than on scenario expectations:

$$\text{order}^* = \arg \max_{o \in \{\text{Simple, Loop, NoBid}\}} R^o_{\text{realised}} \tag{4.13}$$

where $R^{\text{NoBid}}_{\text{realised}} = 0$ by definition. NoBid is therefore selected whenever both Simple's and Loop's realised revenue for the day are negative. This is a distinct criterion from the scenario-based expected revenue $\mathbb{E}[R^o]$ used at submission time to decide *whether* an order is offered in the first place (Simple always clears as a price-taker; Loop's grid-search candidates require $\mathbb{E}[R^o] > 0$ across the 300 scenarios to be submitted at all, Section 4.8.2). The headline win-rate and revenue figures reported throughout Chapter 6 compare what each order type actually earned, not what it expected to earn in advance.

As a secondary risk criterion, the VaR_{10} captures downside exposure over a five-day trading week: a one-in-ten-day adverse outcome corresponds to roughly one adverse day per fortnight, a practically meaningful benchmark for an operator managing rolling cash flows from daily IDA participation. The model separately records the VaR_{10}-optimal order (the type with the highest 10th-percentile revenue across the 300 scenarios) as a secondary risk criterion, reporting how often each order type is VaR_{10}-optimal across the backtest sample. A day-by-day comparison against the realised-revenue criterion above is not developed further in this thesis and is left for future work.

5 | Results

5.1 | Technology Presence: Day-Ahead vs. Intraday

Before examining which order types different technologies rely on, it is worth establishing who participates in each market at all and by which measure.

By distinct participating units, parsing every intraday offer-header file (ICAB) across the full year on disk (February 2025 – February 2026, 1,104 session files) identifies 3,190 units submitting at least one order. Solar PV dominates by unit count: 1,192 distinct units, more than double wind's 423. By offered volume, however, the picture changes substantially: summing every offered MWh across the same year from the intraday offer-detail files (IDET, 1,098 files) shows CCGT alone accounts for 63.1% of offered volume among the five technologies shared with the day-ahead comparison below, against wind's 26.2%, pumped hydro's 9.0%, and negligible shares for nuclear and hydro. Solar PV's volume share (13.4% of all mapped intraday technologies, including those outside the five-category comparison) is an order of magnitude smaller than its unit-count share would suggest: many small solar units each offer comparatively little individually.

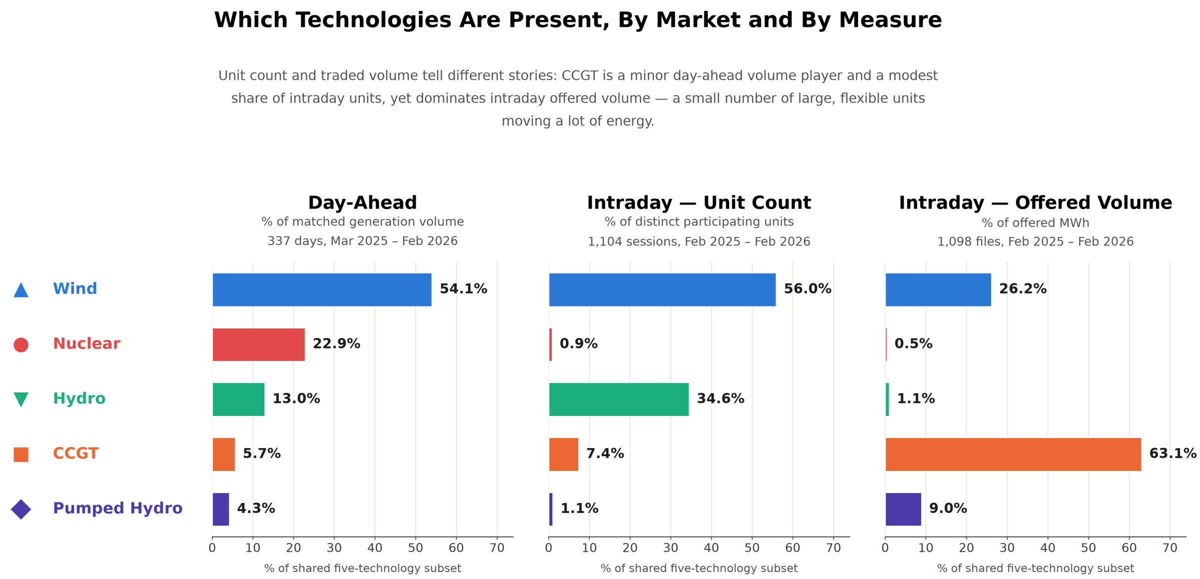

Figure 5.1: Technology presence across OMIE's day-ahead and intraday markets, by three different measures: day-ahead matched generation volume (PDBC, 337 days), intraday distinct participating units (ICAB, 1,104 sessions), and intraday offered volume (IDET, 1,098 files). All percentages are computed within the five technology categories present on all three measures. Source: OMIE (author's processing).

Day-ahead matched volume over the same window shows a third pattern again: wind (54.1%) and nuclear (22.9%) dominate, with CCGT and pumped hydro together represent only 10.0%. Reading the three measures together, a coherent picture emerges. Day-ahead volume reflects installed capacity and baseload commitment, where wind and nuclear are largest. Intraday unit count reflects how many distinct participants show up to trade again after day-ahead closes, where small, weather-dependent solar units are most numerous.

Intraday traded volume reflects how much is actually being moved once those participants are trading, and there CCGT, a comparatively small number of large, flexible thermal units, dominates by a wide margin. Which of these technologies specifically relies on the market's more specialised order types, rather than ordinary simple orders, is examined next.

5.2 | IDA Block Order Usage in Practice

Over the 365-day sample from 1 February 2025 to 31 January 2026, 53 generating units submitted 1,583 block orders to OMIE's three intraday auction sessions. Of these, 572 cleared, giving a 36.1 % acceptance rate. By order count, blocks are 2.25 % of all IDA submissions, but they count for 9.7 % of total scheduled IDA megawatts. This is due to block orders being systematically larger than the simple orders placed alongside them. The instrument appeared on 298 of the 365 trading days (81.6 %), making it a regular feature of IDA order flow.

5.2.1 | Temporal evolution of block order activity

Figure 5.2 plots submissions per day across the full sample. Two things stand out: a step change in activity around August 2025, and an isolated spike in January 2026. From February through July 2025 the series is flat and low, with most days seeing two to four submissions and some days seeing none. From August onward the bars are consistently taller.

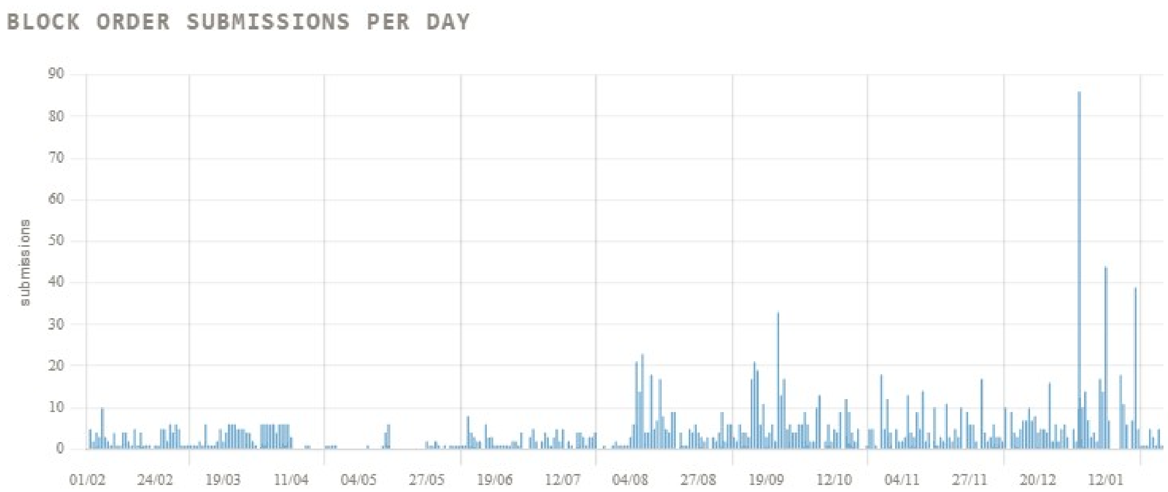

Figure 5.2: Daily block order submissions (IDA1–IDA3 combined), February 2025–January 2026. The dashed vertical line marks the October 2025 transition to quarter-hourly scheduling. Source: OMIE `curva_pibc_uof` (author's processing).

October 2025 is the clearest dividing line, and the most tractable to explain: that month OMIE moved from 24-period hourly to 96-period quarter-hourly scheduling. Before the change, the market averaged 3.30 block submissions per day and blocks appeared on 78.5% of calendar days. After it, the daily average rises to 6.38 and the share of active days rises to 87.8%. The direction of the effect is intuitive. Under hourly scheduling, a generator locking in a four-hour minimum-load block links four periods; the same block in a quarter-hourly world links sixteen. Submitting those sixteen periods as individual simple orders risks partial execution: some slots may clear while others do not, leaving the unit dispatched below its minimum stable generation. A block that executes only when all required periods clear simultaneously removes that risk, so finer-grained scheduling made the instrument relatively more attractive.

However, there is an August jump which precedes October, which means the resolution change consolidated a trend rather than started one. The technology breakdown makes the source identifiable: pumped-hydro units submitted only 6 block orders across the entire February–July period, then 87 in August alone, a fourteen-fold jump in a single month. CCGT submissions also increased, from roughly 50 per month in the first half to 95 in August, but the proportional change for pump-storage is far larger.

Six candidate explanations were checked and ruled out. Correlating monthly submission counts against the monthly average and standard deviation of the marginalpibc auction price rules out a price-driven explanation: the Pearson correlation between monthly submissions and average price is r=0.21r = 0.21 r=0.21, and against price volatility r=0.27r = 0.27 r=0.27, both weak. August's average price (67.4 €/MWh) sits below June (73.1 €/MWh) and December (76.6 €/MWh), months with far fewer submissions, which rules out an unusually high price level as the trigger. Simple-order MW volume for these units shows no coincident ramp around August; trading activity stays flat and noisy through the period, which rules out a general activity increase as the driver. A localized price or weather spike within the window where these units' first block orders clustered (6–9 August) is also ruled out: IDA prices on 15 and 22 July and on 18 August match or exceed the August 6–9 levels, and none of those dates produced a block-order response, so the tight clustering cannot be attributed to an unusual market event. A hydrological trigger, such as reduced pumping activity or tighter execution requirements driven by low reservoir levels, finds no support in the data: MITECO's national water reserve stood at 65.8% of capacity in early August 2025, an unremarkable seasonal figure, and the Duero basin specifically (relevant to unit DUEB) sat 7.4 percentage points above its ten-year average by late August, better than normal rather than worse. A regulatory or system-change trigger is also ruled out: the only relevant OMIE rule change, Instrucción 1/2025, took effect 18 March 2025, introducing quarter-hourly intraday trading and the block-order type itself, nearly five months before the observed clustering. No OMIE circular or CNMC resolution dated in the July–August window addresses block orders or pumped-hydro-specific rules.

New-entrant status rules out a seventh explanation. All seven pumped-hydro units that submitted block orders during this period, including the two whose first block order falls outside the 6–9 August window, had already traded in the intraday market via simple orders since the first week of the sample, between 1 and 18 February 2025: continuous participation, not sporadic entry.

Price level, price volatility, trading volume, a localized market event, hydrological stress, regulatory change, and new-entrant status are all excluded. The tight clustering of first block orders from 5 of the 7 units within a four-day window (6–9 August) points to a shared trigger outside what OMIE's public price, offer, reservoir, and regulatory data can reveal, plausibly an internal commercial or operational decision by the units' owning operators, rather than seven separate operators independently and gradually adopting the instrument. The remaining two units bracket this window loosely: SILB's first block order fell on 19 July, more than two weeks before the cluster, and SLTB's on 26 August, more than two weeks after. The available data cannot identify the specific trigger, and the question stays open.

Table 5.1 shows the month-by-month breakdown.

Table 5.1: Monthly block order submission intensity, February 2025–January 2026

Month	Active days	Total submissions	Avg. per active day
February 2025	26	74	2.8
March 2025	30	96	3.2
April 2025	17	67	3.9
May 2025	9	19	2.1
June 2025	25	53	2.1
July 2025	24	65	2.7
August 2025	29	193	6.7
September 2025	30	231	7.7
October 2025	26	159	6.1
November 2025	26	139	5.3
December 2025	29	168	5.8
January 2026	27	319	11.8

January 2026 stands apart: 11.8 submissions per active day, four times the pre-August baseline. The episode is concentrated in the first two weeks of the month, then drops sharply. The direction breakdown is informative: 66.5 % of January submissions are buy-side (212 buy versus 107 sell), which rules out a thermal minimum-load story and points toward pump-storage charging activity. Why pump-storage operators submitted substantially more buy blocks in January specifically is not clear from the submission data; it may reflect a seasonal pattern in the price dynamics that made charge-side commitments more attractive in that month, but verifying this would require the actual block prices submitted alongside prevailing market prices for those hours.

5.2.2 | Direction and session asymmetry

Of the 1,583 block submissions, 1,010 (63.8 %) are sell-side and 573 (36.2 %) buy-side. The gap in participating units is wider: 49 units submitted sell blocks, against 22 on the buy side. The asymmetry reflects who needs blocks and why. Thermal generators face minimum-load and minimum-operating-time constraints that make partial dispatch either physically impossible or economically irrational; a block that executes only when the full commitment can be met is the natural solution. Storage assets and load-side participants without a minimum-load floor have less structural need for the instrument on the buy side, though the January data show that pump-storage operators do use it there.

Average block duration differs by technology in a way consistent with this reading, though caution is warranted. CCGT units, which place 97.7 % of their block submissions on the sell side, average 8.6 hours per block. Combined-cycle plants typically carry minimum run-time requirements in the range of six to ten hours, and the 8.6-hour average is consistent with the upper end of that interval. The fit is suggestive but has not been verified against individual unit technical specifications. Pump-storage buy blocks average 1.5 hours, suggesting operators are targeting narrow windows rather than committing to extended periods.

Which windows, specifically, can be answered directly from the hours each block spans. Converting each block's $[p_{\text{start}}, p_{\text{end}}]$ period range into clock hours (normalised for the resolution change in October 2025) and counting how often each hour appears gives a clean picture (Figure 5.3). Pump-storage buy blocks concentrate sharply between 12:00 and 16:00, with that four-hour window accounting for 42.4 % of all buy-side hour-coverage versus 27.2 % for the overnight 00:00–06:00 window. This rules out the overnight-trough hypothesis: pump-storage charging blocks target the midday solar price dip, not the low-demand overnight period. CCGT sell blocks peak later, at 17:00–21:00, capturing the evening demand ramp, though the profile is comparatively flat across the day (each six-hour window accounts for 24–26 % of coverage), so the peak-capture behaviour is directional rather than sharply concentrated.

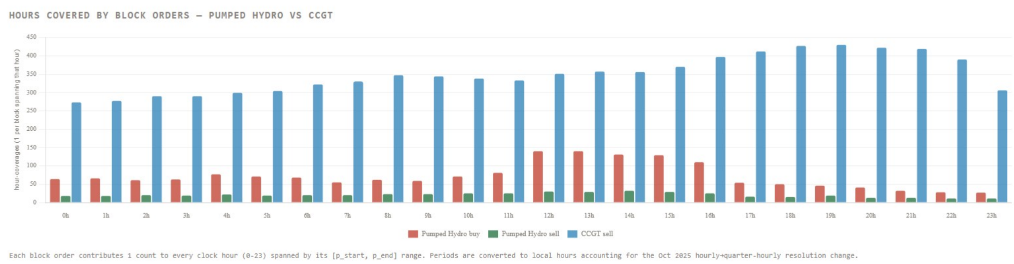

Figure 5.3: Clock hours covered by block orders, by technology and direction. Each block contributes one count to every hour spanned by its declared period range. Pump-storage buy blocks concentrate at 12:00–16:00 (solar midday dip); CCGT sell blocks peak at 17:00–21:00 (evening demand peak). Source: OMIE `curva_pibc_uof` (author's processing).

The near-absence of small (<50 MW) buy-side blocks (37 of 573, or 6.5 %) is easier to interpret directly. A very short charge block delivers minimal energy, and round-trip economics require enough throughput to cover efficiency loss and degradation cost. The threshold below which a block becomes uneconomic depends on individual unit parameters, but the pattern is consistent with this interpretation.

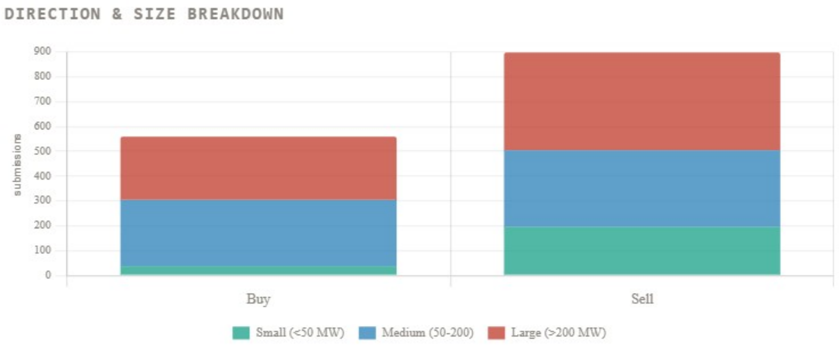

Figure 5.4: Block order submissions by direction and size category. Buy-side activity is concentrated in medium and large orders; small buy blocks account for 6 % of buy-side submissions. Source: OMIE `curva_pibc_uof` (author's processing).

Submission volume falls across sessions: 677 orders (42.8 %) in IDA1, 496 (31.3 %) in IDA2, and 410 (25.9 %) in IDA3 (Figure 5.5). The clearance rate is different; falling from 35.2 % in IDA1 to 31.7 % in IDA 2 and rising to 43.2 % in IDA3 (Figure 5.6). One interpretation of the higher IDA3 rate is a selection effect: participants who have had two sessions to observe conditions submit only high-confidence blocks in the final session, raising the acceptance rate as lower-probability orders drop out. The data do not straightforwardly support this, however: 17 of the 33 units that appear in IDA3 do not appear in IDA1 at all, meaning a substantial fraction of IDA3 participants specifically target the final session rather than selecting out of earlier ones.

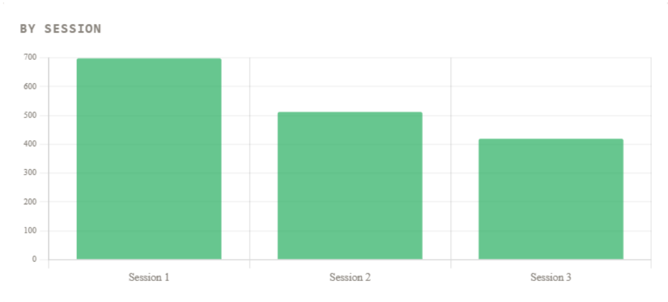

Figure 5.5: Block order submission volume by intraday session. Volume falls from 677 orders in IDA1 to 410 in IDA3. Source: OMIE `curva_pibc_uof` (author's processing).

Figure 5.6: Block order clearance rate by intraday session. IDA3 clears 43.2 % of submissions, eight percentage points above IDA1. Source: OMIE `curva_pibc_uof` (author's processing).

5.2.3 | Technology composition

CCGT units and pump-storage plants account for 87.4 % of all block submissions: 729 orders from 32 CCGT units (46.1 %) and 654 from 15 pumped-hydro units (41.3 %). A single coal unit contributes 116 submissions (7.3 %) and a single gas unit 80 (5.1 %). Four units across nuclear, wind, and conventional hydro submitted 4 orders in total. Table 5.2 shows the breakdown, including the block rate, which is more informative than the raw submission count.

Table 5.2: Block order activity by technology

Technology	Units	Submissions	Share	Block rate[a]	Avg. duration
CCGT	32	729	46.1 %	1.2 %	8.6 h
Pumped Hydro	15	654	41.3 %	2.5 %	1.5 h
Coal	1	116	7.3 %	70.3 %	2.1 h
Gas	1	80	5.1 %	32.4 %	2.2 h
Other	4	4	0.3 %	<1 %	n.a.
Total	**53**	**1,583**			**4.9 h**

[a] Block submissions as a share of total IDA submissions (simple + block) by technology.

Of all IDA offers submitted by CCGT units, only 1.2 % take block form. The 32 units use simple orders for the vast majority of their dispatch, reaching for the block mechanism selectively. The 8.6-hour average duration is consistent with combined-cycle minimum operating times, and the 97.7 % sell orientation is consistent with thermal commitment rather than general trading.

Pump-storage operators use blocks at five times the CCGT rate, and the direction split is informative. Of 654 pumped-hydro block submissions, 547 (83.6 %) are buy-side and 107 (16.4 %) are sell-side. The sell blocks clear at 74.8 %; the buy blocks clear at 28.9 %. Comparing each block's declared price against the `marginalpibc` price realised over its own hours confirms the mechanism directly: pumped-hydro sell blocks are priced, on average, 13.8 €/MWh *below* the market they saw, a safe bid that clears three times out of four. Buy blocks, by contrast, still ask for a 10.4 €/MWh discount below the realised price on average, which the market meets only 28.9 % of the time. The clearance gap between the two sides is therefore a direct consequence of how far each is priced from the market, not an artifact of when the blocks were submitted.

The coal unit (ABO1) submits 70.3 % of its IDA offers as blocks, averaging 2.1 hours each, and clears 9 of 116 (7.8 %). The same price-versus-market comparison explains why: coal's sell blocks are priced 100.6 €/MWh above the market price realised over their own hours on average, roughly eight times the spread CCGT sell blocks carry (12.8 €/MWh, clearing 43.4 % of the time) and far outside anything pumped hydro asks for. This is not evidence of miscalibration in the ordinary sense; a spread this large and this consistent looks like a deliberate standing option at a floor the unit does not expect the market to reach on most days, priced to protect the minimum-generation constraint rather than to trade volume.

Nuclear, wind, and conventional hydro are effectively absent from block order activity, with 4 submissions in total across all three categories, but this is not because these technologies are inactive in the intraday market. Pulling each technology's total IDA offer count (simple plus block, across all 365 days) shows they are structurally absent from the block instrument specifically: solar units submit 1,175,086 IDA offers and zero blocks; wind submits 553,093 offers and 1 block; hydro submits 118,781 offers and zero blocks; nuclear submits 1,853 offers and 2 blocks (0.1 %). Nuclear in the Iberian pool operates as must-run baseload with stable intraday output; wind and solar cannot guarantee multi-period delivery; conventional hydro can ramp flexibly enough to express its operating decisions through simple price-quantity curves. None of the three has a structural reason to demand all-or-nothing execution across multiple periods. The offer counts confirm these units are fully present in IDA, they just never reach for the block mechanism.

The declared minimum prices in block submissions differ sharply by technology. Table 5.3 and Figure 5.7 show the weighted average and the full min–max range for the four technologies with meaningful sample sizes.

Table 5.3: Weighted average minimum price declared in block submissions, by technology

Technology	Weighted avg (€/MWh)	Range (€/MWh)
Gas	176	138–200
Coal	141	100–160
CCGT	108	−30 to 250
Pumped Hydro	44	−15 to 130

Nuclear and wind excluded: fewer than 3 observations each.

Figure 5.7: Required minimum price in block submissions by technology: weighted average with observed min–max range. Source: OMIE `curva_pibc_uof` (author's processing).

The 64 €/MWh gap between CCGT (108 €/MWh) and pumped hydro (44 €/MWh) reflects their underlying cost structures. CCGTs incorporate fuel and start-up costs into their minimum bid prices, whereas pumped hydro can sustain a lower minimum because the expected revenue from selling stored energy exceeds the cost of charging. Five units submitted blocks with negative minimum prices, reaching as low as −30 €/MWh: two CCGT units bidding on both the buy and sell sides, and three pumped-hydro units. For pumped storage, a negative floor on a buy block is economically rational, as operators are willing to pay to consume electricity during periods of surplus if they expect to sell the stored energy later at a higher price. The two CCGTs with negative buy-side floors are less straightforward. At −30 €/MWh, the declared floor represents a threshold that would clear only during an extreme renewable surplus event. In effect, the unit is placing a standing order that says, "If the market falls this far, pay me to consume."

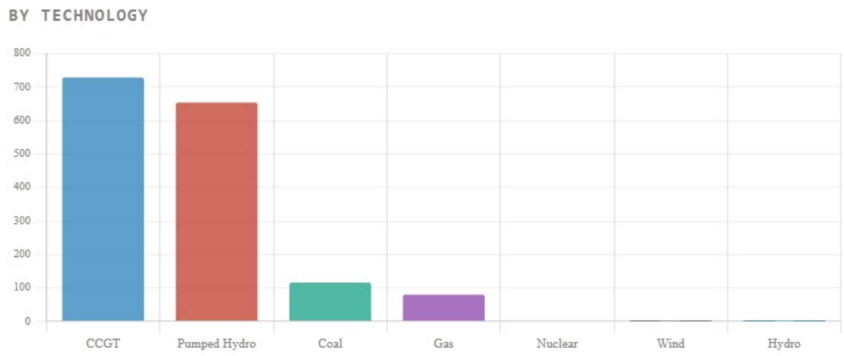

Figure 5.8: Block order submissions by technology. CCGT and pumped hydro account for 88 % of all block orders by count. Source: OMIE `curva_pibc_uof` (author's processing).

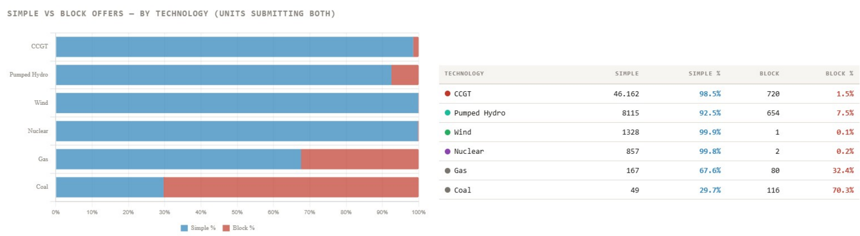

Figure 5.9: Simple versus block share of total IDA submissions by technology, for units submitting both order types. The coal unit allocates 70 % of its IDA offers to blocks; CCGT units allocate 1.2 %. Source: OMIE `curva_pibc_uof` (author's processing).

5.2.4 | Clearance rate and market efficiency

Of the 1,583 block submissions, 572 cleared, giving an overall clearance rate of 36.1 %. Nearly two thirds of block orders are rejected. That figure is less unusual than it might appear: a block order is a limit order applied across multiple periods, and a 36 % acceptance rate may reflect exactly the price floor that rational operators intend to set, accepting that most submissions will not clear. The data cannot tell us whether operators are satisfied with their clearance rates, but the consistency of the pattern across months and technologies suggests it is not the result of systematic miscalibration.

The technology-level rates are direct observations. CCGT clears 43.5 % and pump-storage 36.4 %, close to the overall figure. Within pump-storage, sell blocks clear at 74.8 % while buy blocks clear at 28.9 %. The coal unit clears 7.8 % (9 of 116 submissions) and gas 8.8 % (7 of 80). These rates are shown in Figure 5.10.

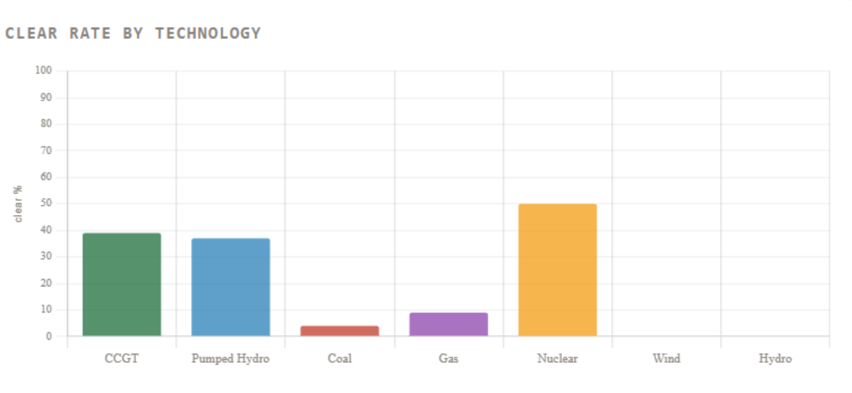

Figure 5.10: Block order clearance rate by technology. Coal and gas units clear fewer than 9 % of submissions; CCGT and pumped hydro clear approximately 38 %. Source: OMIE `curva_pibc_uof` (author's processing).

Across units, price level and clearance rate move in opposite directions. Gas (176 €/MWh average) and coal (141 €/MWh) clear below 9 % of their submissions. Pump-storage sell blocks, averaging around 44 €/MWh, clear at 72 %. The highest unit-level clear rates in the dataset, 52–55 %, come from a small group of CCGT units that bid both buy and sell sides at a weighted average of around 88 €/MWh. Units priced well above the expected clearing range use blocks to signal a minimum-run commitment rather than to trade volume; those priced competitively clear at substantially higher rates. Figure 5.11 shows this relationship across all 53 units, with bubble size proportional to submission count.

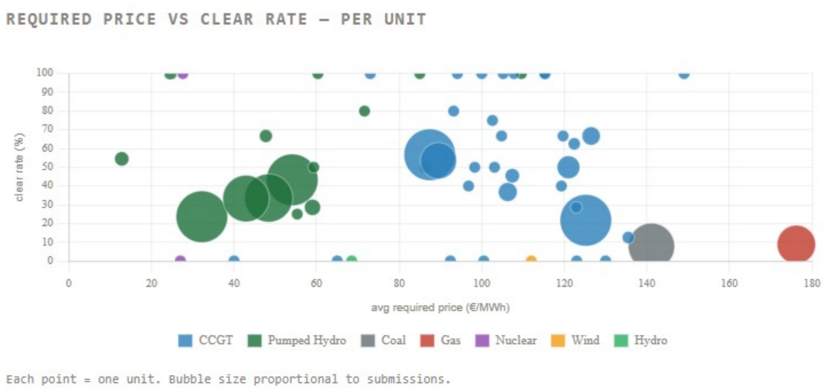

Figure 5.11: Required minimum price versus clearance rate, per unit. Bubble size is proportional to total block submissions. Source: OMIE `curva_pibc_uof` (author's processing).

Because the declared minimum price alone conflates a unit's absolute pricing strategy with the price level of the hours it targets, it is more informative to measure how far each block was priced from the market it saw: the average `marginalpibc` price over the block's own $[p_{\text{start}}, p_{\text{end}}]$ span (Figure 5.12). This spread predicts clearance almost mechanically: coal sell blocks average 100.6 €/MWh above their own realised market price and clear 0.9 % of the time; gas sell blocks average 29.3 €/MWh above and clear 8.8 %; CCGT sell blocks average 12.8 €/MWh above and clear 43.4 %; pumped-hydro sell blocks average 13.8 €/MWh *below* the realised market and clear 74.8 %. The ordering is monotonic across all four technology/direction pairs with a meaningful sample, and it holds regardless of the absolute price level each technology operates at; what matters is the distance from the market the block faced.

Figure 5.12: Declared block price versus the **`marginalpibc`** price realised over the block's own hours, per block. Left: average spread by technology (positive = priced further from clearing). Right: required price against realised market price, coloured by cleared/not cleared; the diagonal marks price = market. Source: OMIE **`curva_pibc_uof`** (author's processing).

Figure 5.13 shows that the high-volume periods, August–September 2025 and January 2026, do not produce proportionally more accepted orders. The gap between submitted and cleared orders widens when volume rises, meaning the additional submissions during those periods clear at a below-average rate. Given that August's average market price is unremarkable (Section 5.2.1), the widening gap more likely reflects the surge bringing in submissions priced further from the market, in line with the spread-clearance relationship above, not a general rise in market prices during those months.

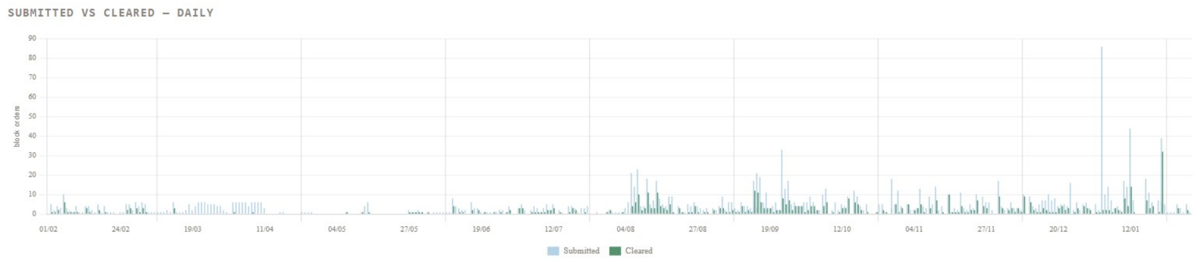

Figure 5.13: Daily block order submissions versus accepted orders, February 2025–January 2026. The gap between the two series widens in high-volume periods. Source: OMIE **`curva_pibc_uof`** (author's processing).

The relationship between elapsed time in the sample and clearance rate is also worth noting. Monthly clear rate correlates moderately with month index across the sample ($r = 0.51$), rising from an unstable 3–12% in the first quarter to a comparatively stable 40–46% from May 2025 onward. A learning-curve explanation is plausible: as operators gain experience with the instrument, they may calibrate the declared floor closer to the price the block is likely to realise, improving clearance without changing the underlying pricing logic.

5.2.5 | Seasonal composition of block-order activity

The technology breakdown in Table 5.2 and the August 2025 surge discussed in Section 5.2.1 both concern a single month or a static annual total. A broader question sits behind both: across the full sample, do CCGT and pumped hydro follow different seasonal patterns in how they use the block instrument at all? Grouping the same monthly data underlying Table 5.1 into meteorological seasons[1] gives a cleaner picture than the August episode alone, though it does not resolve that episode.

Table 5.4 reports unit-days (days where at least one block has been sent) and distinct active units by season. CCGT participation is broadly stable across winter, summer, and autumn, 27–33 % of the year's unit-days and 20–22 distinct units in each, with spring the clear outlier at 9.8 % of unit-days and only 5 distinct units. Pumped hydro is far more concentrated: essentially absent in spring (1 unit-day out of 192 for the full year), rising through summer (15.6 %, 9 units) to an autumn peak of 45.8 %, the largest of any season and the most distinct units active in a single season (13) for this technology.

Table 5.4: Unit-days and distinct active units by season, CCGT and Pumped Hydro

	CCGT			Pumped Hydro		
Season	**Unit-days**	**Share**	**Units**	**Unit-days**	**Share**	**Units**
Winter	202	32.5 %	21	73	38.0 %	11
Spring	61	9.8 %	5	1	0.5 %	1
Summer	193	31.0 %	22	30	15.6 %	9
Autumn	166	26.7 %	20	88	45.8 %	13
Total	**622**		**32**	**192**		**15**

Table 5.5 restates the same seasons in MW offered, MWh cleared, and clear rate. CCGT's clear rate collapses in spring, 16.1 %, against 42–53 % in every other season, and the collapse holds across two of the three months rather than one: March clears 3 of 38 submissions (7.9 %) and April 1 of 16 (6.2 %), while May's 6 of 8 (75.0 %) rests on too small a base to offset the pattern. Section 5.2.4 shows that clearance across the full sample tracks how far a block is priced from the market it saw, almost mechanically; the spring collapse is consistent with that same mechanism if CCGT blocks in March and April were priced further from realised prices than in other months, which would follow from thinner thermal dispatch margins when hydro and renewable output are seasonally highest. That price-versus-market spread has not been broken out by season here, so this remains a plausible reading rather than a verified one. Pumped hydro shows a different shape: its clear rate stays comparatively stable across the seasons it participates in at all, 28.2 % winter, 44.6 % summer, 43.3 % autumn, so its seasonal story is almost entirely about whether units submit at all rather than how well those submissions are priced once they do.

[1] Winter = December, January, February; Spring = March, April, May; Summer = June, July, August; Autumn = September, October, November. Because the sample runs from 1 February 2025 to 31 January 2026, "Winter" here combines a partial four-week February 2025 fragment with a full December 2025–January 2026 pair, not a complete three-month season on both ends.

Table 5.5: MW offered, MWh cleared, and clear rate by season, CCGT and Pumped Hydro

	CCGT			Pumped Hydro		
Season	**MW offered**	**MWh cleared**	**Clear rate**	**MW offered**	**MWh cleared**	**Clear rate**
Winter	40,288	179,111	43.1 %	51,493	23,025	28.2 %
Spring	8,157	2,964	16.1 %	65	33	100.0 %*
Summer	50,916	167,863	42.0 %	11,249	4,399	44.6 %
Autumn	51,650	130,971	53.3 %	39,402	21,346	43.3 %
Total	**151,011**	**480,909**	**43.5 %**	**102,209**	**48,803**	**36.4 %**

*Pumped hydro spring $n = 1$; see footnote above.

The volatility of intraday prices gives a direct mechanism for both findings above. Sampling the same backtest's raw IDA Session-1 price series across eight dates per month for the full year shows that spring carries the lowest intraday price volatility of any season, on both measures: an average daily price standard deviation of 22.08 €/MWh and an average daily price range of 81.06 €/MWh, essentially tied with winter (23.10 €/MWh and 81.39 €/MWh) and roughly 35–40 % below summer (35.67 €/MWh, 111.13 €/MWh) and autumn (34.44 €/MWh, 117.53 €/MWh) on both metrics. A block order only pays off if the price spread it captures, whether between its own periods or against the marginal clearing price, is large enough to justify the all-or-nothing clearing risk described in Section 5.2.4. When volatility is this structurally compressed, both the incentive to submit a block order and the probability that a submitted block clears fall together. The most likely explanation for the spring pattern, both CCGT's collapse in participation and its collapse in clear rate, is this compression in the underlying price spread, not a separate cause for each.

This seasonal pattern bears on the August 2025 surge without explaining it. The five newly-adopting pumped-hydro units clustering within 6–9 August (Section 5.2.1) sit at the boundary between summer's 15.6 % share of unit-days and autumn's 45.8 % peak: block-order activity for this technology was already trending upward into that window and continued rising afterward, consistent with a broader seasonal expansion of the instrument's use across the second half of the sample. That context does not identify why the specific four-day window in early August saw five separate operators' individual adoption decisions cluster so tightly; the six candidate explanations ruled out in Section 5.2.1 still leave that clustering unexplained. The seasonal and August-specific findings are complementary rather than substitutable: one describes a gradual six-month shift in participation, the other an unresolved four-day coincidence inside it.

5.2.6 | From block orders to battery storage: a gap in the market

The above results reveal a structural pattern: block orders in the IDA are the instrument of technologies with physical indivisibility constraints. CCGTs need a guaranteed minimum run window to recover start-up costs; pumped hydro needs sustained pumping periods to fill the reservoir efficiently. Both have a physical reason to demand all-or-nothing execution. A battery storage system has neither constraint. It can ramp in seconds, maintain any dispatch duration, and accept partial fills without loss of efficiency. The all-or-nothing character of a block order therefore provides no operational benefit to a BESS, while introducing non-clearing risk on every session where the price ends up below the declared floor.

Scale adds a second, independent objection. Figure 5.14 shows average block size by technology. CCGT blocks average 240 MW per order; pump-storage blocks range from 70 to 220 MW across units. A BESS rated at 5 MW falls an order of magnitude below even the smallest pump-storage commitment. The problem is not just design logic but magnitude: the instrument is not sized for a small storage asset, regardless of how the minimum-price condition is calibrated.

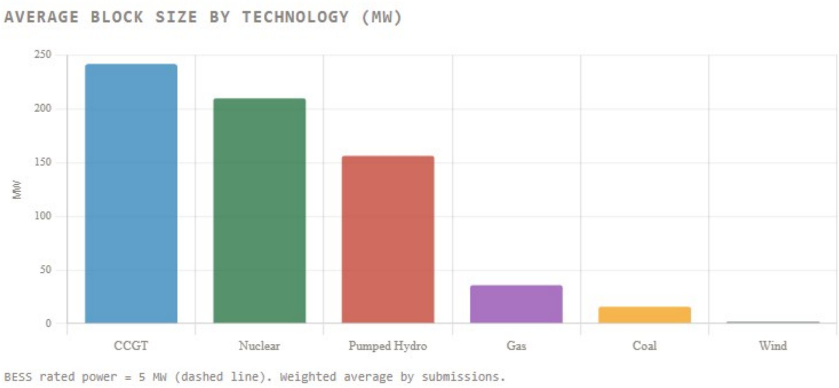

Figure 5.14: Average block size by technology (MW). The dashed line marks the 5 MW rated power of the BESS modelled in this thesis. CCGT and pumped-hydro commitments are one to two orders of magnitude larger. Source: OMIE `curva_pibc_uof` (author's processing).

Even for well-calibrated thermal units, the overall block clearance rate is only 43.5 % for CCGT and 7.8 % for coal. A BESS bidding with a price forecast subject to forecasting error, particularly underestimation of price spikes, would face the same calibration difficulty without the physical run-time rationale that makes the risk worth taking for a gas turbine.

Block orders, then, do not fit batteries. Yet Spain's PNIEC 2023–2030 targets 22.5 GW of installed energy storage by 2030, with BESS making up a growing share of that fleet. As batteries become significant participants in OMIE's intraday sessions, the question of which order types serve their operational needs becomes pressing. One candidate, which is not available at the moment in OMIE's intraday rulebook, is the **Loop C88 order**: a paired buy–sell structure that encodes a minimum net revenue condition across the full arbitrage cycle rather than a minimum price on one leg. Unlike a block order, it is designed for round-trip energy assets. The following sections evaluate how a 10 MWh / 5 MW BESS performs in the day-ahead market under simple, block, and Loop C88 order strategies, and how well each captures the available price spread.

Summary

Block orders account for 2.25 % of all IDA offers by count but 9.7 % of cleared megawatt-hours, reflecting their large per-order energy commitment. CCGT blocks average 240 MW per order; pump-storage blocks range from 70 to 220 MW across units. These are not instruments designed for small assets.

Among generating units, **CCGTs** are the primary users on the sell side: they use blocks to guarantee a minimum run window (average 8.6 h) at a weighted average minimum price of 108 €/MWh, priced 12.8 €/MWh above the market their blocks see, achieving a 44 % clear rate; their sell blocks concentrate at 17:00–21:00, capturing the evening demand ramp. **Pumped hydro** units use blocks in both directions, averaging 44 €/MWh across submissions: sell blocks for discharge are priced 13.8 €/MWh below their realised market and clear at 75 %, while buy blocks for pumping still ask for a 10.4 €/MWh discount and clear at only 29 %; buy-block hours concentrate at 12:00–16:00, targeting the midday solar price dip rather than the overnight demand trough. **Thermal baseload** (coal, gas) quotes prices of 141 and 176 €/MWh respectively, priced 100.6 and 29.3 €/MWh above their realised market on average, and clears below 9 % of submissions, making their block activity economically marginal on a per-order basis. **Nuclear, wind, and solar** are not absent from the intraday market; together they submit over 1.8 million IDA offers across the sample, but are structurally absent from the block instrument specifically (combined block share < 0.1 %), consistent with having no minimum-run or minimum-load constraint that a block would help satisfy.

Across the dataset, how far a block is priced from the market it saw predicts clearance almost mechanically, and the relationship is monotonic across every technology/direction pair with a meaningful sample: units pricing near or below the realised market clear; units pricing well above it rarely do. The August–September 2025 surge in submissions is a step change in behaviour, not a response to price conditions ($r = 0.21$ between monthly submissions and average price). For a BESS, the instrument fails on two independent grounds: the all-or-nothing structure offers no operational benefit to a fully flexible asset, and the typical block size is an order of magnitude above the battery's rated power.

5.3 | Intraday BESS Backtest Results

5.3.1 | Backtest Design and Overview

Before turning to the intraday market, the same battery model, LEAR forecaster, and order-type framework (Section 4.8) were first run on the OMIE day-ahead auction as a validation cross-check, confirming the pipeline behaved sensibly on the better-documented day-ahead session before it was applied to the intraday auctions (Section 5.3.13). The intraday auctions are the focus of this chapter; the day-ahead run is referred to only occasionally below, as an incidental point of comparison rather than a result in its own right.

The intraday backtest evaluates the LEAR-based LP optimiser on the OMIE IDA auctions (Sessions 1, 2, and 3) over 433 days of available price data within a window spanning 5 February 2025 to 11 May 2026; a small number of calendar days within this window (28 of 461) lack a usable OMIE price file and are excluded from the sample. The battery modelled is a 10 MWh / 5 MW system (a two-hour battery) with round-trip efficiency $\eta = 0.92 \times 0.92 = 84.6$ %, a state-of-charge (SoC) operating range of 5–95 %, and a degradation cost of 10 €/MWh discharged. These hardware parameters are identical to those used in the day-ahead model.

A central design choice is to treat each IDA session as an independent optimisation problem, equivalent to running the same 10 MWh / 5 MW battery model three times, once per session, each with its own self-contained energy budget. At the start of every session the battery is reset to $S_0 = S_{\min} = 5\%$, and the LP enforces a terminal SoC-neutrality constraint ($S_{\text{end}} = S_0$), so each session's revenue is evaluated in isolation. There is no SoC carry-over between sessions, nor from any day-ahead schedule into Session 1. Session 1 and Session 2 each span the full 24-hour day (periods P1–P96, 96 quarter-hour intervals); Session 3 covers only the second half of the day (periods P49–P96, 48 intervals), reflecting the structural design of the IDA3 auction schedule. This isolates the standalone arbitrage value each session offers on its own price information, independent of what any other session or the day-ahead market contributes.

Forecasting uses the LEAR model specified in Section 4.5: 96 independent LASSO regressions, one per quarter-hour, trained on an expanding window of prior IDA prices for each session, consistent with the LEAR-LASSO framework described in Section 4.5 (the training window grows by one day with each session rather than discarding older observations). Dispatch uncertainty is represented through 300 Monte Carlo price scenarios per session drawn from $\mathcal{N}(\hat{p}_t, \hat{\sigma}_t)$, where $\hat{\sigma}_t$ is the 14-day rolling standard deviation of forecast errors. Two order types are computed independently for every session: **Simple** (LP dispatch, price-taker, always clears) and **Loop C88** (a paired charge-plus-discharge block, accepted jointly only if the realised price spread between legs covers the declared degradation cost at actual clearing prices). Two independent methods were evaluated for scheduling Loop's charge/discharge windows – a grid search over the LEAR point forecast maximising expected revenue $\mathbb{E}[R]$ across the 300 Monte Carlo scenarios, and a fast rule-based heuristic that targets full-duration capacity utilisation directly from the LEAR forecast without Monte Carlo scenarios – and the results reported throughout this chapter use the rule-based method, which Section 5.3.2 shows achieves substantially higher realised Loop revenue in every session; the grid-search method is retained in that section as the rejected alternative, for comparison. Both orders' realised revenues are then compared directly: for each session the better-performing order at actual prices is reported, and for each day the order type with the higher total realised revenue summed across sessions is taken as the day's outcome. An oracle benchmark solves the LP on the actual, realised IDA prices for each session, bounding the maximum achievable revenue under perfect foresight.

1,561 €	**1,952 €**	**18.6 %**	**24**
avg daily revenue (Sessions 1+2+3, Simple)	avg oracle revenue (perfect foresight)	avg gap to oracle (best order, S1+S2+S3)	loss days (combined S1+S2+S3, Simple actual rev < 0)

Simple orders achieve the higher realised revenue in the large majority of session-days across all three sessions: 91.2 % in Session 1, 84.8 % in Session 2, and 82.2 % in Session 3. Loop C88, scheduled with the rule-based method (Section 5.3.2), wins on 5.0 %, 7.6 %, and 4.6 % of session-days respectively; the remainder in each session (3.8 %, 7.6 %, and 13.2 %) are classified as no bid.

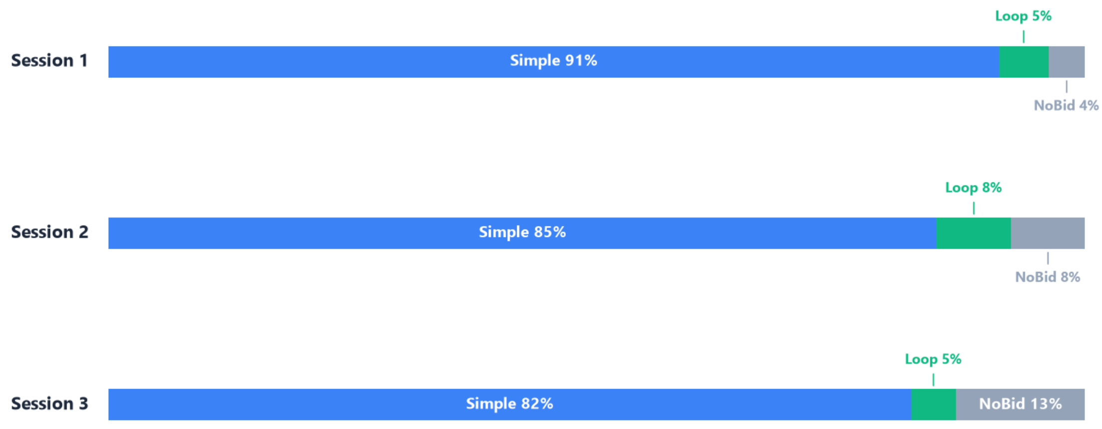

Figure 5.15: Order type win share by session. Session 1: 398 post-burn-in days; Session 2: 395; Session 3: 393.

Block orders were initially modelled, but were never selected as a primary instrument, consistent with the IDA block order finding presented earlier in this chapter (Section 5.2): a battery has no ramp-time constraint and no minimum stable generation level, so it will rather opt to tie its blocks or bid in each quarter-hour independently via simple offers. The remainder of this section thus analyses Simple and Loop performance in depth.

5.3.2 | Two Approaches to Loop Scheduling: Grid Search versus a Rule-Based Heuristic

Simple's dispatch schedule has 96 independent quarter-hourly decisions per session, so a full LP solve is the natural choice: the problem is exactly the kind of large, structured linear program an LP solver is built for, and Section 5.3 shows Simple's flexibility is itself the main driver of its advantage over Loop. Loop's own scheduling problem is much smaller – one or two charge/discharge window pairs per session, not 96 independent periods – and two genuinely different methods for solving it were implemented and backtested against the full 433-day sample.

Grid search (originally reported). The first method exhaustively scores candidate charge/discharge window pairs drawn from a fixed menu of block lengths (2, 4, 8, 12, 16, and 24 quarter-hour periods), ranking candidates by their expected revenue $\mathbb{E}[R]$ across 300 Monte Carlo price scenarios, and to represent the battery's two-cycle capacity it repeats the search once more in the periods left over after the first loop is chosen.

Rule-based heuristic (used throughout this chapter). The second method drops the Monte Carlo step entirely and works directly from the LEAR point forecast: for each of several candidate charge durations, biased toward the battery's full per-cycle capacity, it finds the cheapest available window of that length, pairs it with the best-matching discharge window sized to exactly what round-trip efficiency allows that charge to sell back, and scores the resulting pair by its forecast net profit (revenue minus degradation cost) rather than by average price spread. This directly targets capacity utilisation rather than the single sharpest average spread, motivated by Modo Energy's (2023) industry finding that "maximum revenues are available for a two-hour system completing 2 cycles per day" in intraday battery arbitrage, and by Ferrara et al.'s (2024) result that a simple, forecast-driven rule-based scheduler "is able to closely approach the performance achieved through LP in all the tested scenarios," with the two revenue curves in their comparison "visually indistinguishable." Ferrara et al.'s result was obtained for a full LP-scale scheduling problem; Loop's is a much smaller, two-decision problem, so the case for a lightweight heuristic replacing an exhaustive search is, if anything, more direct here than in their original setting.

Table 5.6: Grid search versus rule-based Loop scheduling, full backtest (433 days, Sessions 1–3). Per-session rows use each session's own post-burn-in day count (398, 395, 393 respectively, as in Table 5.7); the Combined row uses the smaller 387-day intersection on which all three sessions are simultaneously post-burn-in (Table 5.8), so it is not the arithmetic sum of the three rows above it.

Session	Grid search total (€)	Rule-based total (€)	Improvement	Rule-based clears
Session 1	53,704	157,841	+193.9 %	95.2 %
Session 2	61,064	143,899	+135.7 %	89.1 %
Session 3	59,287	117,706	+98.5 %	82.2 %
Combined (387-day intersection)	173,474	416,969	+140.4 %	—

The rule-based method more than doubles Loop's total realised revenue in every session and realises a genuinely profitable outcome (both legs clear and the realised spread covers degradation) on 82–95 % of days, against under 2 % for the grid search. This is a stricter standard than clearing alone: the rule-based method's forecast-driven filter proposes a pairing on essentially every session-day, but a proposed pairing that clears can still realise a loss when the actual prices diverge from the forecast that justified submitting it (Section 5.3.11); the 82–95 % figure counts only the subset that clears *and* turns out profitable.

The mechanism is directly visible in how much of the battery's capacity each method actually uses: across all legs found by the rule-based search, the average charge window is 6.6 periods (1.6 hours) and the average discharge window is 4.7 periods (1.2 hours) – the shorter discharge duration follows directly from the round-trip efficiency constraint, which caps discharge energy at $\eta = 0.846$ times the charge energy for the same rated power, not from the search settling for a narrow window the way the grid search's spread-ranked candidates did. The grid search, by contrast, systematically converges on the shortest legal block at each step (typically 4 charge periods paired with 2 discharge periods, Section 5.3.14) because ranking candidates by average price spread rewards concentrating on the single sharpest instant in a trough or peak, with no term in the objective that rewards trading more energy at a still-profitable price.

The two methods also differ in how they enforce state-of-charge feasibility for a second cycle. The grid search's second-loop search excludes only the exact periods the first loop occupies, without requiring the first loop's discharge to complete before a second loop begins charging;

checking the running state of charge implied by its reported two-cycle candidates shows this can require the battery to hold approximately 10 MWh simultaneously at some point in the session – above the 9 MWh usable window (5–95 % of the 10 MWh nameplate) used throughout this model. The rule-based method's two-cycle search instead considers a second cycle confined either entirely before the first cycle's charge window or entirely after its discharge window, selecting whichever span yields the better forecast profit, so the battery is never asked to hold both cycles' energy simultaneously regardless of which cycle occurs first in time. This also lets the search recover a genuine earlier-in-day opportunity that an ordering restricted to always searching forward from the first-chosen cycle would miss even when no physical conflict exists; Section 5.3.8 gives a concrete example. The results in this chapter use only physically feasible schedules throughout: the grid-search totals in Table 5.6 are computed under the same state-of-charge feasibility requirement as the rule-based method, so the two columns are directly comparable.

5.3.3 | Operational days of the model

LEAR burn-in period. The LEAR model requires a rolling burn-in window of approximately 35 days of prior session prices before it can fit the 96 LASSO models. The first 35 trading days of the sample (5 February to approximately 14 March 2025) fall within this initialisation window for Sessions 1, 2, and 3 alike; on these days the model consistently returns no bid and actual revenue is zero. All per-day failure-mode analysis and monthly charts in the sections below are computed on the post-burn-in sample only, to avoid attributing structural no-bid outcomes to model failure.

Day-count reconciliation. Table 5.7 reports 398 operative days for Session 1, 395 for Session 2, and 393 for Session 3, rather than the 430 (Session 1), 427 (Session 2), and 428 (Session 3) days for which price data are available within the window described above. The difference in each case (32–35 days) is exactly the burn-in period described above; these days are excluded from all headline figures because the model has no forecast to act on, not because of a data gap. The combined Sessions 1+2+3 sample of 387 days is the intersection of days on which all three sessions have completed their burn-in.

5.3.4 | Aggregated Results

Table 5.7 reports per-session results.

Table 5.7: Per-session backtest results (Section 5.3.3).

Session	Days	Avg Simple (€/day)	Avg Loop (€/day)	Avg Oracle (€/day)	Gap
Session 1 (P1–P96)	398	+584.3	+396.6	+665.0	11.6 %
Session 2 (P1–P96)	395	+533.2	+364.3	+709.7	23.2 %
Session 3 (P49–P96)	393	+420.7	+299.5	+555.7	22.0 %
Sessions 1+2+3 (combined)	387	+1,560.8	+1,077.4	+1,951.9	**18.6 %**

The combined Sessions 1+2+3 headline is presented in Table 5.8. Over 387 operative days (the subset on which all three sessions have simultaneously completed burn-in), Simple orders generated total revenue of 604,020 € against an oracle of 755,389 €. At the day level (summing each session's revenue before comparing order types, matching Table 5.8's own convention), Simple wins on 350 of the 387 combined days (90.4 %), Loop on 15 days (3.9 %), and the remaining 22 days (5.7 %) are classified as no bid. The best-order actual revenue captures **81.4 %** of the per-session oracle value, with an average gap of 18.6 %. Session 2 has the widest point-estimate gap (23.2 %), followed closely by Session 3 (22.0 %) and, well separated from both, Session 1 (11.6 %). Session 1's gap is visibly separated from Session 2's and Session 3's due to using the day-ahead price as an anchor, while Session 2 and Session 3 sit close enough together that their point-estimate ordering should be read cautiously rather than as a demonstrated ranking; formal significance testing of these differences is left for future work.

Table 5.8: Combined Sessions 1+2+3 headline metrics (387 days, post burn-in, within 5 February 2025 – 11 May 2026).

Metric	Value
Avg Simple revenue / day	+1,560.8 €
Total Simple (387 days)	+604,020 €
Avg Loop revenue / day	+1,077.4 €
Total Loop (387 days)	+416,969 €
Avg oracle revenue / day	+1,951.9 €
Total oracle (387 days)	+755,389 €
% of oracle value captured (best order)	81.4 %
Avg gap to oracle (best order)	18.6 %

5.3.5 | Revenue Distribution and Seasonal Pattern

Revenue generated by the best available order type each session (Simple, Loop, or NoBid) forms a wide distribution, driven by the heterogeneity of IDA price spreads across sessions and seasons. Table 5.9 shows the daily revenue breakdown by bucket, combining all three sessions' contributions on each of the 387 days on which all three have completed burn-in.

Table 5.9: Daily revenue distribution under the best-available-order strategy (post burn-in, Sessions 1+2+3 combined, 387 days — the same day set used in Table 5.8).

Revenue range	Days	Share	Avg oracle (€)
Near-zero (0–100 €)	20	5.2 %	117
Low (100–500 €)	46	11.9 %	604
Medium (500–1,500 €)	109	28.2 %	1,412
High (>1,500 €)	212	54.8 %	2,695
Total (post burn-in)	387	100 %	—

More than half of all active days (54.8 %) generate over 1,500 € in combined-session revenue, a pattern driven by Spain's pronounced solar and wind-driven intraday price spreads. No combined day falls below 0 €, since each session's own contribution to the total is never allowed to go below zero: a session where neither Simple nor Loop is profitable is simply not traded (NoBid), which floors that session's contribution at exactly zero rather than a loss. The 20 near-zero days (0–100 €) with an oracle well above that range represent cases where spread existed in at least one session but the model failed to capture it; these are analysed in Section 5.3.7.

The seasonal pattern, shown in Figure 5.16 for all three sessions individually, is pronounced and consistent across sessions. Summer months (June–September 2025) average 515–868 €/day across the three sessions, driven by solar duck-curve dynamics that produce deep midday price troughs and sharp evening demand peaks. Winter months (December 2025, January 2026) fall to 105–336 €/day as renewable generation profiles flatten and intraday spreads compress. February 2026 is an outlier in every session: Session 1 falls to 170 €/day, Session 3 to 31 €/day, and Session 2 turns *negative*, averaging −17 €/day, the only negative monthly average anywhere in the sample – concentrated in a cluster of near-zero or negative IDA clearing prices caused by a sustained renewable surplus event (Section 5.3.7). The pattern's consistency across sessions, despite each session drawing on an independently trained LEAR forecaster, indicates the underlying driver is the market's own seasonal price structure rather than an artifact of any single session's forecasting pipeline.

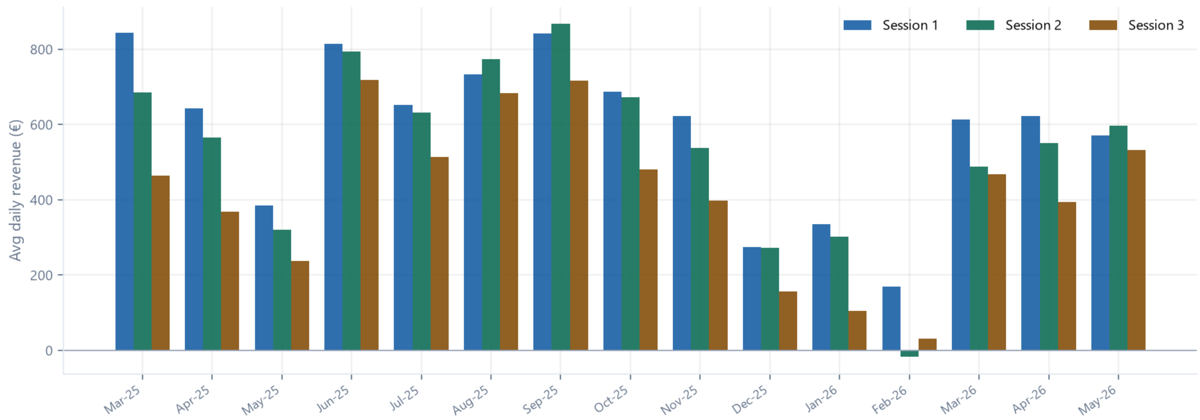

Figure 5.16: Average daily revenue by month, Sessions 1, 2, and 3 (Simple strategy). February 2026 reflects a sustained near-zero/negative IDA price event driven by renewable surplus (Section 5.3.7), visible in all three sessions and pushing Session 2's monthly average below zero. March 2025 includes only the post-burn-in days.

Figure 5.17 shows the same information at daily resolution across the full post-burn-in sample, combining all three sessions' revenue against their combined perfect-foresight oracle. The February 2026 trough is visible as a sustained run of near-zero and negative bars, distinct in character from the day-to-day volatility that surrounds it elsewhere in the sample. Figure 5.18 shows the distribution of the daily gap to oracle underlying each session's own average gap (Table 5.7), faceted by session since the three distributions differ meaningfully in shape: Session 1's gap is strongly right-skewed and concentrated near zero (178 of 398 days, 45 %, within 50 € of the oracle; only 35 days, 9 %, with a gap of 150 € or more), while Session 2's distribution is shifted markedly wider (only 46 of 395 days, 12 %, within 50 €; 161 days, 41 %, at 150 € or more) and Session 3 sits between the two (106 of 393 days, 27 %, within 50 €; 107 days, 27 %, at 150 € or more). This, together with the day-ahead anchor, directly explains why Session 1's average gap to oracle (11.6 %) is substantially narrower than Session 2's (23.2 %) and Session 3's (22.0 %): the difference is not driven by a handful of extreme days in any one session, but partly by a systematically wider spread of everyday gaps in Sessions 2 and 3.

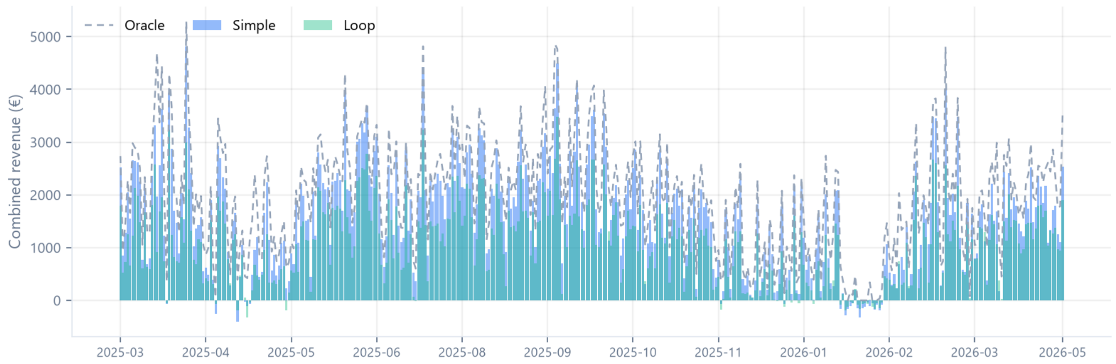

Figure 5.17: Daily revenue, Sessions 1+2+3 combined, 387 post-burn-in days. Simple (blue) and Loop (green) bars against the perfect-foresight combined oracle (grey dashed line). The February 2026 renewable-surplus episode appears as a sustained band of near-zero and negative bars around 2026-02.

The headline averages in Table 5.8 describe a typical day, but a single mean doesn't show the consistency of what the battery actually earns: an average of 1,560.8 €/day could describe a market where every day looks alike, or one where a handful of exceptional days carry the whole sample while most days earn little. The distinction matters directly for the failure-mode analysis that follows in Section 5.3.7, since it first requires separating genuinely low-opportunity days (where the oracle itself is small and a low realised revenue is simply correct) from days where real opportunity existed but the model failed to capture it. Table 5.9 bins each of the 387 post-burn-in days by combined realised revenue and reports the average oracle alongside each bucket, so that the two cases can be told apart before any individual day is examined in detail.

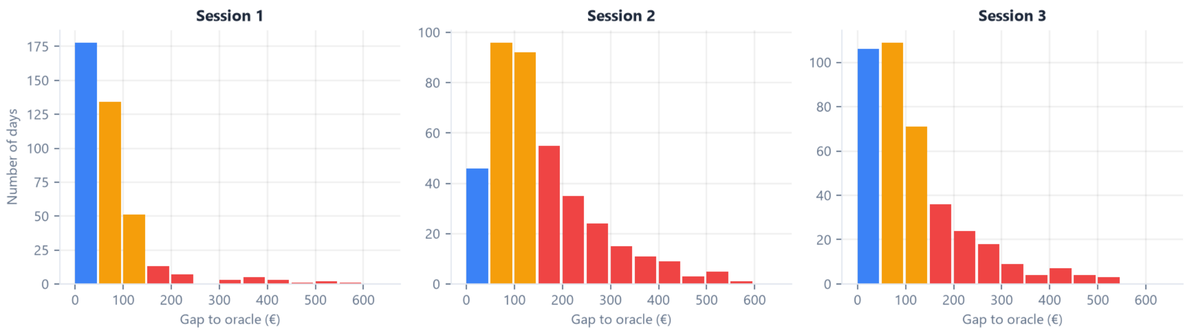

Figure 5.18: Distribution of the daily gap to oracle (best order), by session. Bars are coloured by gap magnitude (blue <50 €, amber 50–150 €, red ≥150 €). Session 1's distribution is markedly more concentrated near zero than Session 2's or Session 3's.

5.3.6 | Simple's Own Operating Profile

Before turning to the Simple-versus-Loop comparison that occupies the rest of this chapter, it is worth isolating what the Simple order does on its own terms, independent of any comparison against Loop or the oracle. The LP formulation places no restriction on how many times the battery may charge or discharge within a session beyond the physical state-of-charge and power limits (Section 4.6): a schedule with one charge block and one discharge block is neither favoured nor penalised relative to a schedule with six. In practice, the optimiser exploits that freedom whenever the forecast price curve rewards it, and the resulting dispatch pattern is the clearest illustration of what a period-by-period order type buys a battery that a paired-block order structurally cannot.

In the following sections, different figures will be shown. These are present in the html connected to the code which stores and represents visually the behaviour of the model for each day of the sample. A small manual explaining how to read each figure can be found in the Appendix A.1.

Figure 5.19 shows Session 2 on 17 September 2025, a day on which Simple captures 1,665.7 € against an oracle of 1,751.4 € – a gap of only 85.7 € (4.9 %) – while Loop, evaluated on the same forecast and the same realised prices, reaches just 1,281.5 €, 23.1 % less than Simple. The mechanism is not that Simple slices one broad ramp into finer pieces; it is that the realised price path oscillates sharply over intervals of thirty to sixty minutes – most visibly between 08:00 and 10:00, where the price spikes, drops, spikes again, and drops a second time inside two hours, and again between 17:00 and 19:00 – and the LP chases each of those short-lived swings individually, producing eight distinct charge/discharge pulses across the session rather than one or two sustained blocks. Loop's clearing condition settles a leg against a single declared price limit for the whole span it covers; asked to represent a price path that reverses direction four times in two hours, the best it can do is average across the swing and capture the net difference between one high point and one low point, discarding whichever intermediate reversals fall inside its own window. This is a capability, not a tuning problem: no choice of window width or price limit lets a two-leg order re-enter a position it has already closed within the same session, so the intermediate reversals are unreachable for Loop by construction, not merely under-exploited by this particular heuristic.

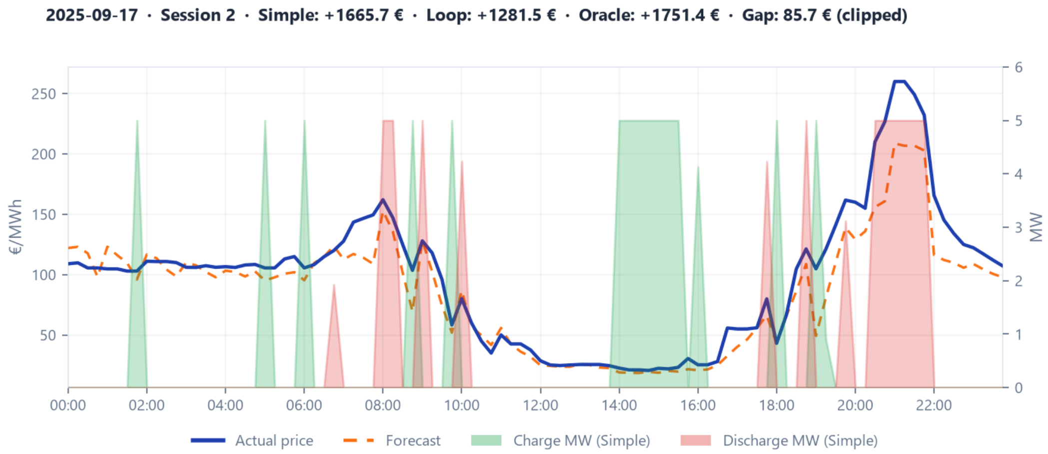

Figure 5.19: 17 September 2025, Session 2. Simple (1,665.7 €) closes to within 4.9 % of the 1,751.4 € oracle by fragmenting its schedule into eight separate charge/discharge pulses, each chasing a short-lived reversal in the realised price – most visibly the four direction changes packed into the 08:00–10:00 window alone. Loop, on the same forecast and prices, reaches only 1,281.5 € (23.1 % less than Simple): its single charge-discharge pairing can settle against one price limit per leg, so it can capture at most the net swing across a window, never the intermediate reversals inside it.

This day is rather the rule than the exception in the near-oracle sample in Table 5.10 it draws from. Across that table, Simple's per-session revenue advantage over Loop on the very same forecast and the very same realised prices ranges from roughly 29 % to 195 %, and the mechanism is the same fragmentation visible here: Loop is not failing to find the best two-leg pairing available to it, it is being asked to summarise a price curve that Simple does not need to summarise at all. Section 5.3.12 returns to this same fragmentation statistic – 98.7 % of session-days, in order to maximize their revenue, require more than the two legs a single Loop pairing can express.

The biggest advantage of the simple orders is that, if the price forecasting model is good enough, on a large share of active days the model achieves revenue close to the oracle. Table 5.10 presents a selection of days where the gap is within 4 % and combined revenue exceeds 1,500 €. On these days, LEAR correctly identifies the dominant intraday price structure, typically a midday solar trough followed by an afternoon demand peak, and the LP allocates charge to the valley and discharge to the peak across all three active sessions.

Table 5.10: Sample of simple order, near-oracle days: combined revenue within 4 % of oracle (Sessions 1+2+3 combined, post burn-in).

Date	Simple (€)	Oracle (€)	Gap	Loop (€)
25 Jun 2025	3,687	3,748	1.6 %	2,776
23 Jun 2025	3,355	3,442	2.5 %	2,502
27 Jun 2025	2,159	2,216	2.6 %	1,508
09 Jan 2026	2,248	2,321	3.2 %	1,293
14 Jun 2025	2,282	2,359	3.3 %	1,735
21 Jun 2025	3,003	3,109	3.4 %	2,268
25 Sep 2025	4,041	4,200	3.8 %	2,660

A consistent pattern emerges across all seven days in Table 5.10 and across all three sessions within each one: Loop generates substantial realised revenue in every single session (1,293–2,776 € combined across the three sessions) but Simple still wins every session on every one of these days. This is not a close contest decided by a handful of periods: Simple's per-session advantage over Loop ranges from roughly 29 % to as much as 195 % of Loop's own revenue on these particular days. The Simple LP captures more by optimising each quarter-hour independently across all three sessions simultaneously, rather than accepting Loop's paired-legs clearance constraint in even one of them; near-oracle outcomes on these days come from LEAR correctly identifying the cheap and expensive periods in every session, not from any single session dominating the total.

OMIE's uniform marginal-price settlement plays a significant role on days where LEAR underestimates the spread magnitude. A battery that submits a sell offer at 50 €/MWh into a session that clears at 180 €/MWh receives 180 €/MWh on all dispatched energy, provided only that its declared price falls below the session's marginal price. Correct ordering of cheap and expensive hours is therefore sufficient for near-oracle outcomes, even when the forecast underestimates absolute price levels.

Figures 5.20 and 5.21 show the dispatch chart for two of the days in Table 5.10, and illustrate a feature the summary statistics alone do not show: the Simple LP does not necessarily dispatch as a single charge block and a single discharge block per session. It freely fragments its schedule into as many separate windows as the forecast price curve has distinct troughs and peaks, which is precisely what drives the gap between Simple and Loop performance on these two days.

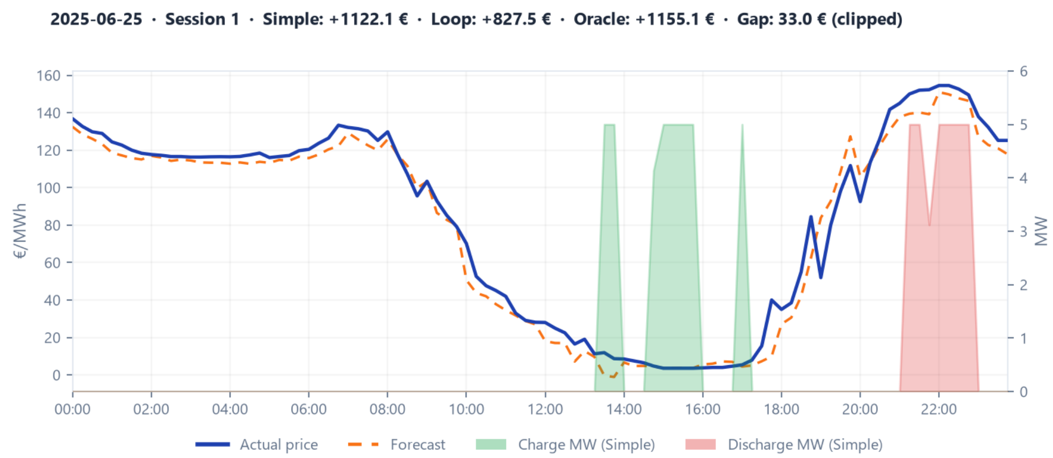

Figure 5.20: 25 June 2025, Session 1 (1,122 € of the 3,687 € combined Sessions 1+2+3 total in Table 5.10). Rather than a single charge block spanning the 13:00–17:00 trough, the LP splits its charging into *three* short pulses (13:00, 15:00, and 17:00) that track the noisy real-time fluctuations within the trough more tightly than one continuous block could. Discharge similarly splits into two blocks around the 21:00 peak, with a brief pause near 22:00 where the realised price dips before recovering. This fragmentation is the concrete reason Simple (1,122 €) outperforms Loop (827 €) on this day: even with the two-cycle search, Loop's paired blocks must average over contiguous periods rather than exploit each narrow fluctuation individually the way the period-by-period LP does.

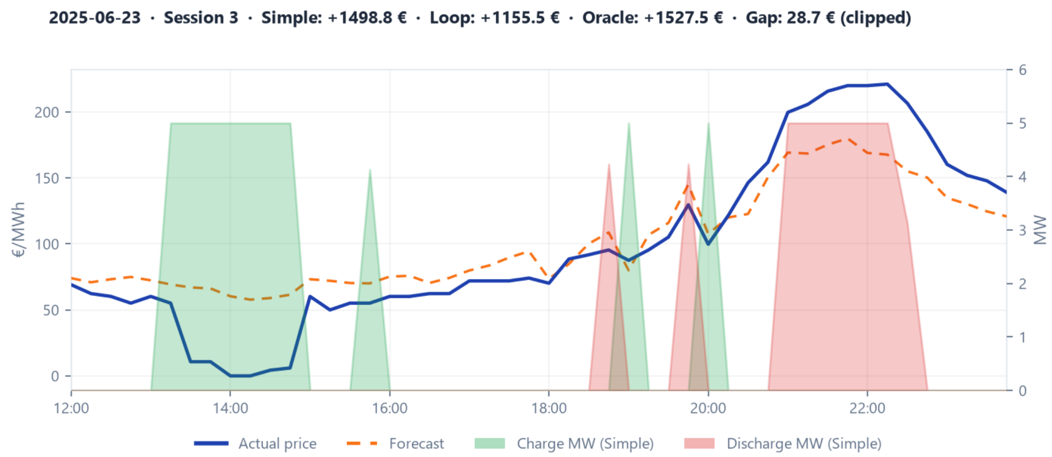

Figure 5.21: 23 June 2025, Session 3 (P49–P96; part of the same 3,355 € near-oracle combined day in Table 5.10). Beyond the single deep midday trough (13:00–15:00) captured by one continuous charge block, the LP also opportunistically trades a second, smaller ramp in the early evening: three short alternating charge/discharge pulses between 18:30 and 20:00, ahead of the main evening discharge block at 21:00–23:00. This secondary trade is only visible in the dispatch chart, not in the session's aggregate revenue figures, and shows the Simple LP exploiting more structure in the price curve than a single charge-discharge pair could, in a session whose day-length window (P49–P96) is structurally different from Sessions 1 and 2's full 24-hour span.

5.3.7 | Failure Cluster and Failure Mode Taxonomy

Session 1 has 18 Simple loss days (post burn-in), Session 2 has 32, and Session 3 has 48. None is uniformly distributed across the sample, but the three sessions differ sharply in *how* their losses cluster. Session 1's losses are heavily concentrated in a single window: 13 of 18 fall within the first three weeks of February 2026, with a further cluster in late January 2026, and 17 of the 18 have an oracle below 50 €, meaning almost none of Session 1's losses reflect a real missed opportunity. Sessions 2 and 3 tell a different story. February 2026 accounts for a smaller share of their losses (19 of 32, 59 %, for Session 2; 10 of 48, 21 %, for Session 3), and roughly half their loss days have a substantially positive oracle (17 of 32 for Session 2, 26 of 48 for Session 3) — days where a real, meaningful arbitrage opportunity existed and the model's own pre-session expected revenue was also large and positive, confirming the forecast was confidently wrong about the market delivering that spread, not merely dispatching into a market that genuinely offered nothing. Figure 5.23 illustrates a representative Session 1 loss day from the February cluster.

The February 2026 cluster has a single root cause, visible across all three sessions. IDA clearing prices during 1–18 February 2026 averaged below 1 €/MWh on most days in Session 1, with Session 2's monthly average turning outright negative and Session 3 falling to 31 €/day. This reflects a prolonged renewable surplus event in the Spanish electricity system: high wind and solar output combined with limited export capacity suppressed intraday prices to near-zero or negative levels for an extended period, market-wide and session-independent. The battery lost money not because the forecast was wrong about the spread direction but because the LEAR model, trained independently for each session, dispatched based on a forecast that still anticipated a normal-spread day in every session:

a sustained, novel regime of this kind is a small minority of the expanding training history at this point in the sample, and the 14-day rolling window used for the dispersion estimate $\hat{\sigma}_t$ likewise had no recent observations of this regime to draw on in any of the three independently-trained forecasters. The resulting throughput paid degradation costs against a market that offered near-nothing in return.

The loss days across the full post-burn-in sample, all three sessions combined, fall into two structural categories.

Type I: Phantom spread (81 days across the three sessions: 17 in Session 1, 15 in Session 2, 22 in Session 3, plus a further 27 positive-oracle days in Sessions 2 and 3 that share the identical mechanism at larger scale). The model forecasts a spread and dispatches aggressively, but the actual market does not deliver it: revenue is negative because degradation cost is paid on energy throughput that generates no net margin, and in every case the model's own pre-session expected revenue was substantially positive, confirming the failure is a forecast-shape error rather than a timing or direction error. Representative examples spanning all three sessions: Session 1, 31 January 2026 (E[simple]=+188 €, actual=−93 €, oracle=36 €); Session 2, 24 March 2026 (E[simple]=+1,150 €, actual=−0.2 €, oracle=894 €), the largest single instance of this failure mode in the sample; Session 3, 26 October 2025 (E[simple]=+24 €, actual=−313 €, oracle=448 €).

Type III: Wrong dispatch window (a small residual in each session). On a minority of loss days the oracle is substantially positive and the model's own pre-session expectation is comparatively modest rather than confidently large, indicating the model located the wrong charge or discharge window rather than mis-forecasting the existence of a spread altogether. Session 1's clearest example is 30 April 2025: the model selects the wrong charge window, charges into elevated actual prices, and discharges at a price only marginally higher, leaving a spread too thin to cover degradation costs (actual =−41 €, E[simple]=+338 €, oracle=140 €). This failure type, which recurs in the day-ahead cross-check, is rarer in the intraday data than Type I: the shorter IDA forecast horizon reduces the probability that LEAR mislocates the cheap-hour window relative to a full 24-hour day-ahead forecast. Figure 5.22 shows the dispatch chart for this day.

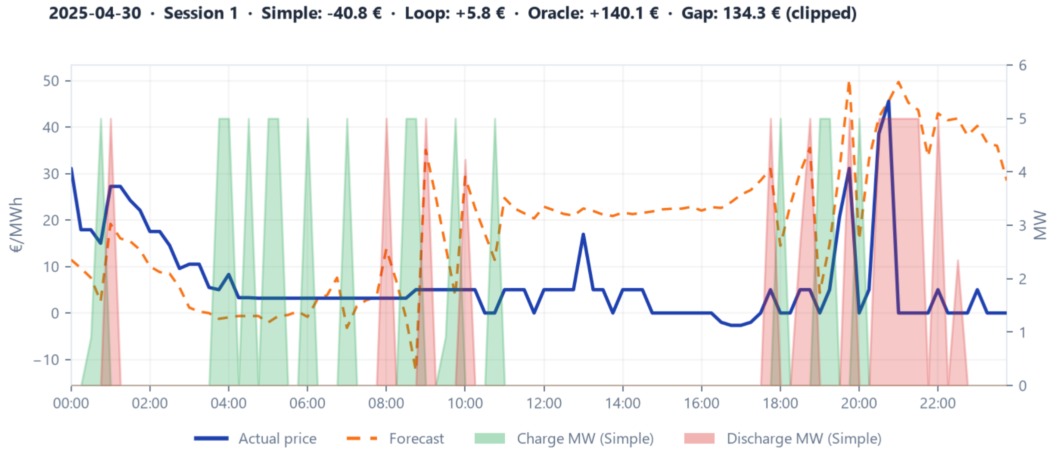

Figure 5.22: 30 April 2025, Session 1: Type III failure. The model splits charge across several morning and midday windows and discharges into a late-afternoon rise that only marginally exceeds the charge price, leaving a spread too thin to cover the degradation cost once summed across the dispatched energy.

Table 5.11: Representative loss days by failure type, spanning all three sessions (post burn-in).

Date	Ses.	Type	Actual (€)	Oracle (€)	E[Simple] (€)
31 Jan 2026	1	I	−93	36	+188
06 Feb 2026	1	I	−92	0	+129
01 Mar 2026	1	I	−84	24	+220
24 Mar 2026	2	I	−0.2	894	+1,150
25 Feb 2026	2	I	−58	638	+564
26 Oct 2025	3	I	−313	448	+24
11 Apr 2026	3	I	−126	485	+199
30 Apr 2025	1	III	−41	140	+338

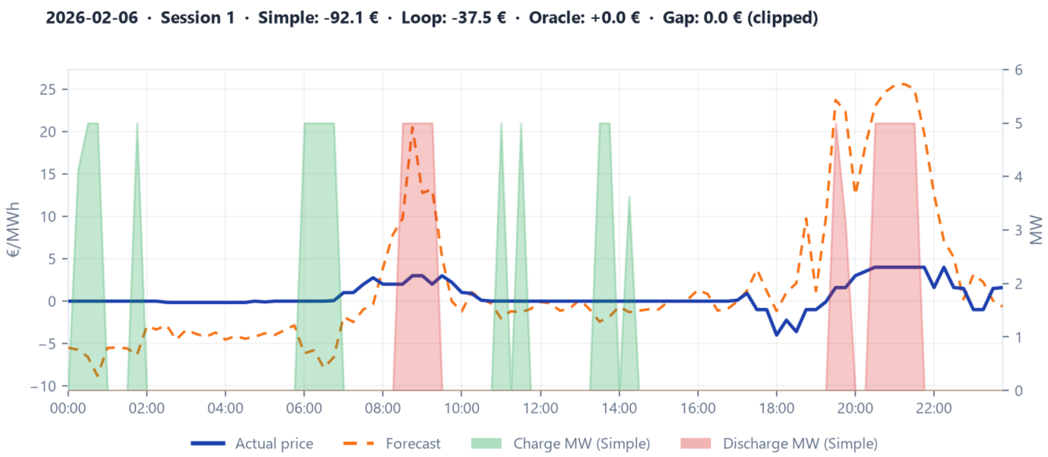

Figure 5.23: 6 February 2026, Session 1: Type I failure from the February 2026 cluster. The LEAR forecast (orange dashed) predicts a normal-spread day and the model dispatches aggressively. The actual IDA clearing price (blue) is near zero throughout, offering no spread. The battery pays degradation costs on the dispatched energy and recovers nothing from the flat market. The oracle on this day is also close to zero. Figure 5.26 shows the same 6 February episode striking Session 2 at an identical forecast-spike location, confirming the event was market-wide.

A notable difference from the day-ahead cross-check is the *absence* of Type II failures (missed spikes: oracle large, model bids nothing) in the post-burn-in IDA sample, in any of the three sessions. In the day-ahead run, 27 October 2025 and 31 October 2025 produced oracle revenues of 950 € and 465 € respectively while the model submitted no bids, because the 14-day residual window contained no precedent for those extreme price events. In the IDA sessions, the shorter forecast horizon and the expanding LEAR training window, updated with each successive session, appear sufficient to prevent these categorical missed-spike events across all three independently-trained forecasters: when large spreads appear in the IDA, the model's recent price history is close enough in time to identify them, at least partially.

5.3.8 | Loop Order Analysis: Selection, Clearance, and Revenue

Using the rule-based scheduling method (Section 5.3.2), Loop C88 achieves higher realised revenue than Simple on 20 of 398 Session 1 days (5.0 %), 30 of 395 Session 2 days (7.6 %), and 18 of 393 Session 3 days (4.6 %) — 68 session-level wins across 1,186 session-days (5.7 %), corresponding to 15 of 387 combined days (3.9 %) once each day's three sessions are summed before comparison. This is a substantially higher win share than the grid-search method achieves on the same sample (9 session-level wins, 0.8 %, Section 5.3.2), consistent with the far higher realised revenue and clearing rate documented there. Every genuine Loop win involved an order that cleared the C88 net-margin condition $\lambda_{h_2} - \lambda_{h_1} \geq c_{\text{deg}}/\eta$ at the actual clearing prices, and every leg respects the round-trip efficiency constraint (discharge energy capped at $\eta = 0.846$ times charge energy) and the real-time state-of-charge ordering described in Section 5.3.2; no session-day in the sample shows Loop's reported revenue exceeding the oracle's. Loop's own realised revenue can be negative when a submitted pairing clears but fails to cover degradation at actual prices (Section 5.3.11); these session-days are never counted as Loop wins here, since a win requires Loop's realised revenue to be both positive and higher than Simple's.

Figure 5.24 shows Simple versus Loop realised revenue across all 1,186 active session-days, pooling all three sessions and colouring each point by session. The three sessions occupy essentially the same region of the plot, with no session systematically offset from the others: a substantial cluster of points across all three sits above the Simple = Loop line, concentrated where both strategies earn positive but moderate revenue rather than at the extremes, confirming that Loop's wins are driven by capturing a single genuine, full-duration ramp cleanly, not by exploiting days where Simple's own forecast-driven schedule fails outright, and that this mechanism is common to all three sessions rather than a Session-1-specific pattern. The cluster of points below zero on both axes, visible for all three sessions, is Loop's own realised losses (Section 5.3.11): unlike Simple, whose losses are floored only by the realised market, Loop's losses arise specifically when a cleared pairing fails to cover degradation.

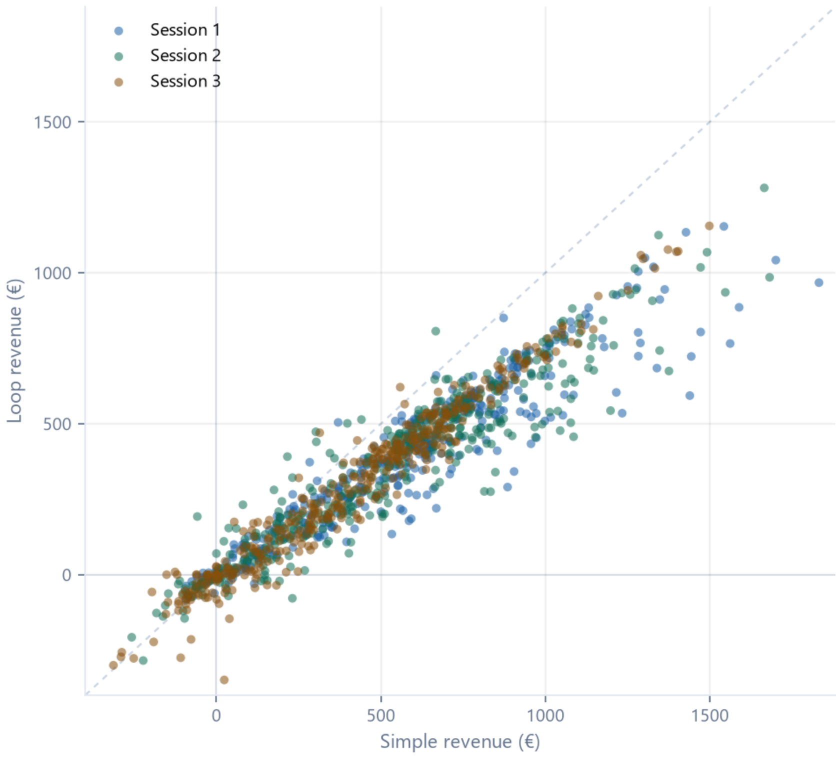

Figure 5.24: Simple vs. Loop realised revenue, all three sessions pooled (1,186 active session-days). Point colour follows session (Session 1, 2, or 3); the dashed line marks Simple = Loop.

Monthly distribution. With 68 session-level Loop wins spread across most months of the sample but concentrated in the shoulder-season and winter months ("What the winning days have in common", below), Loop's advantage tracks whichever individual days happen to have one strong, sustained, typically solar-driven ramp that Simple's forecast-driven schedule prices slightly less efficiently than the rule-based Loop's full-duration capture of the same ramp. A full month-by-month seasonal characterisation comparable to the Session 1 failure-mode analysis in Section 5.3.7 has not been carried out for Loop wins specifically and is left for future work.

5.3.9 | Loop Day-Level Analysis

Table 5.12 lists a representative selection of the 68 session-level days across the three sessions where the rule-based Loop search out-earned Simple, spanning the strongest wins, the weakest (narrowest-margin) wins, and days where Loop avoided a Simple loss outright.

Table 5.12: Representative days where Loop C88, scheduled with the rule-based method, achieved higher realised revenue than Simple. Oracle is the perfect-foresight Simple LP for the same session, charged the battery's $\eta = 0.846$ round-trip efficiency loss; settled on declared MW and price per leg.

Date	Ses.	Loop (€)	Simple (€)	Oracle (€)	Note
30 Oct 2025	1	372.6	283.5	799.2	Two-cycle win
09 Oct 2025	2	391.1	215.6	721.1	Strong single-cycle win
09 Oct 2025	3	470.2	314.1	764.2	Strong single-cycle win
25 Feb 2026	2	192.6	−58.2	637.7	Loop avoids a Simple loss
11 Apr 2026	3	9.1	−125.9	485.0	Loop avoids a large Simple loss
18 Apr 2025	3	27.3	−0.6	65.0	Loop avoids a Simple loss
04 Mar 2026	2	806.8	666.8	1,242.9	Largest absolute Loop win
10 May 2025	2	90.2	89.6	343.1	Weakest positive margin
17 Feb 2026	1	87.7	85.0	139.3	Weakest positive margin

Two features stand out, both different from the grid-search method's day-level pattern (Section 5.3.2). First, 59 of the 68 session-level wins (86.8 %) use only a single charge/discharge cycle, so most of Loop's advantage comes from pricing one full-duration cycle substantially better than the grid search's short, spread-maximising legs rather than from trading twice; the remaining 9 wins (13.2 %) genuinely use two cycles, made possible by the rule-based search considering a second cycle confined either entirely before or entirely after the first cycle's window rather than only ever searching forward from wherever the first cycle happens to fall (Section 5.3.2). A second cycle is a second independent forecast bet layered onto the first, and it is disproportionately represented among Loop's own loss days as well as its wins (Section 5.3.11): when the second leg's forecast is wrong, its loss is summed directly into the day's total, which is part of why two-cycle days make up a smaller share of Loop's wins here than of its overall submissions. 30 October 2025 in Session 1 illustrates the two-cycle mechanism directly: the first cycle (13:30–20:00, capturing the session's dominant midday-to-evening ramp) is chosen first because it is individually more profitable, and the second cycle is then found entirely *before* it, in an early-morning window (04:15–09:30) that a forward-only search could never reach once the first cycle is fixed; the second leg's own realised profit was modest but still positive here (372.6 € total against 799.2 € oracle), unlike the second-cycle losses discussed in Section 5.3.11. Second, wins are broadly distributed across all three sessions (20 in Session 1, 30 in Session 2, 18 in Session 3), roughly in proportion to each session's own operative-day count, rather than concentrated in whichever session Simple performs weakest in.

What the winning days have in common. Pooling all 68 session-level wins reveals a consistent price signature, and it is not the one a spike-chasing intuition would predict. The median realised intraday price spread on a Loop-win day is 49.1 €/MWh, against

94.6 €/MWh on the remaining 1,118 active session-days: winning days sit, on average, only around the 27th percentile of the full sample's spread distribution. Loop wins on *calmer*, not more volatile, days. This follows directly from the instrument's own structure: because a single Loop leg is settled on the *average* price over its declared charge and discharge windows rather than the peak instantaneous price either window contains, a large but narrow spike is diluted away as soon as the declared window extends even a few periods beyond it, whereas a smooth, single-direction ramp loses almost nothing to that averaging. Consistent with this, Simple's own dispatch is, if anything, slightly *more* fragmented on Loop-win days than on the rest of the sample (median 12 blocks per session-day, against 10 for all active session-days): the price curve on these days still has some structure Simple exploits with several small legs, but none of that structure is a spike sharp enough to reward a short, concentrated block the way the (rejected) grid-search method's ranking favoured.

The timing of the charge and discharge windows on winning days is also systematic. Charge windows cluster in two bands – overnight (00:00–04:00, 26 of 77 legs) and midday (12:00–16:00, 26 of 77 legs) – consistent with the solar duck-curve dynamics documented throughout this sample (Section 5.3.7), where midday PV output drives prices toward zero; discharge windows cluster overwhelmingly in the evening ramp (20:00–24:00, 32 of 77 legs), with a secondary cluster at the morning peak (08:00–12:00, 17 of 77 legs).

The same solar signature appears seasonally: 67 of the 68 wins (98.5 %) fall in February–May or October–January, the shoulder-season and winter months in which Spain's PV output most often collapses midday prices to near zero, while June through September contribute a single win across the entire sample. This is not because summer prices are calmer: Session 1's median daily price spread is actually highest in June–September (110.5 €/MWh) of any four-month block in the sample, exceeding the shoulder/winter months' median (81.9 €/MWh), and the shape of that spread is not systematically spikier either, so a narrow-spike explanation for summer's near-absence from the win list does not hold up against the data.

Simple's own oracle-capture rate is, if anything, slightly higher in summer (median 91.8 %) than in the rest of the sample (median 88.0 %), consistent with Simple's period-by-period flexibility capturing summer's larger but still fragmented price swings at least as efficiently as it captures the shoulder-season's smaller ones, leaving Loop with comparatively less room to beat it; establishing the precise mechanism behind this seasonal pattern is left for future work. At the monthly level, a general association between flatter prices and a higher Loop win rate does hold, though only moderately: comparing each calendar month's median Session 1 price spread against that month's Loop win rate across all sessions gives a Pearson correlation of -0.45, and February 2026, the single flattest month in the sample (median spread 23.8 €/MWh), also has the highest win rate (16.0 % of session-days). The relationship is not tight enough to treat flatness alone as a reliable predictor, however: March 2025 has a high median spread (118.0 €/MWh) and a below-average win rate, while June–September's own high spreads coincide with a near-zero win rate, as noted above. February and March 2026 alone account for 24 wins (35.3 % of the total), consistent with the same February 2026 low-demand, high-renewable regime shift documented as a forecast failure mode in Section 5.3.7; since that episode drives both the month's unusually flat prices and its unusually high win rate simultaneously, the correlation above may partly reflect this one event rather than a general flatness effect independent of it.

On 11 April 2026, 18 April 2025, and 25 February 2026, Loop wins not by earning a large positive revenue but by avoiding a Simple loss: Simple loses 125.9 €, 0.6 €, and 58.2 € respectively on these three days while Loop earns a modest 9.1 €, 27.3 €, and a genuinely substantial 192.6 €. These are examples of the risk-reduction property discussed in Section 5.3.11 below, though that section also shows this pattern does not hold on every Simple loss day: Loop's own realised revenue is negative on 15 of Session 1's 18 loss days. At the other end of the distribution, the largest single-day Loop win in the sample, 4 March 2026 in Session 2 (Loop 806.8 vs. Simple 666.8 €), shows the rule-based method capturing a large, genuine ramp cleanly (Figure 5.27) while still leaving a wide gap to the 1,242.9 € oracle, underscoring that even Loop's best days capture only a fraction of the available spread relative to a perfect-foresight schedule.

5.3.10 | Session-by-Session Validation: Simple and Loop's Dispatch Mechanism

The dispatch-level mechanism behind Simple's fragmentation advantage and Loop's occasional wins, illustrated above for Session 1, holds with the same character in Sessions 2 and 3. Figures 5.25–5.29 show one near-oracle day, one Type I failure day, and one Loop-win day from each of Session 2 and Session 3.

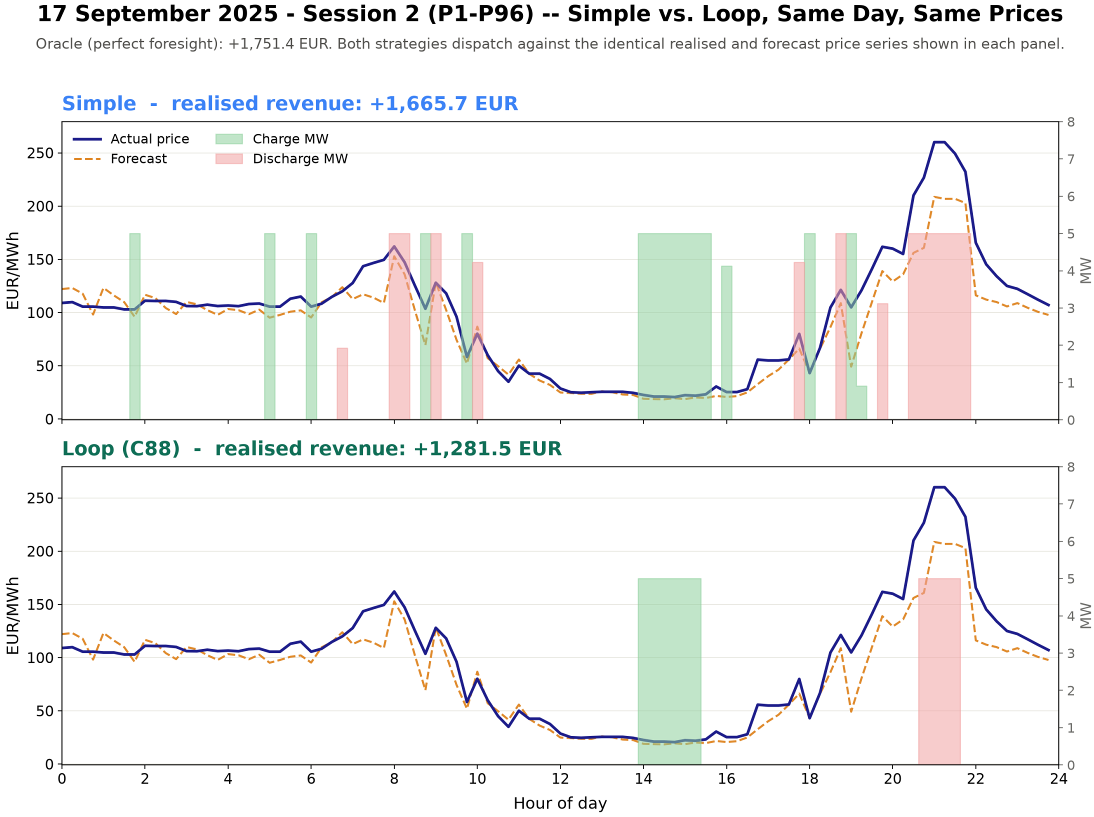

Figure 5.25: 17 September 2025, Session 2: Simple (top) against Loop (bottom) on the identical realised and forecast price series. The realised price has only one genuine trough (00:00–07:00), yet the LP still fragments its charge allocation into several separate short pulses rather than one continuous block, and adds a brief, isolated discharge blip near 07:00 before the main discharge block at 08:00–10:00 (top panel). As in Figure 5.20, this fine-grained fragmentation is not visible in the summary revenue figures alone, and it is why Simple (1,665.7 €) outperforms Loop (1,281.5 €) here even though Loop still finds and clears a genuinely profitable single-cycle pairing (bottom panel): the price structure offers more separate opportunities than one paired block can capture at once.

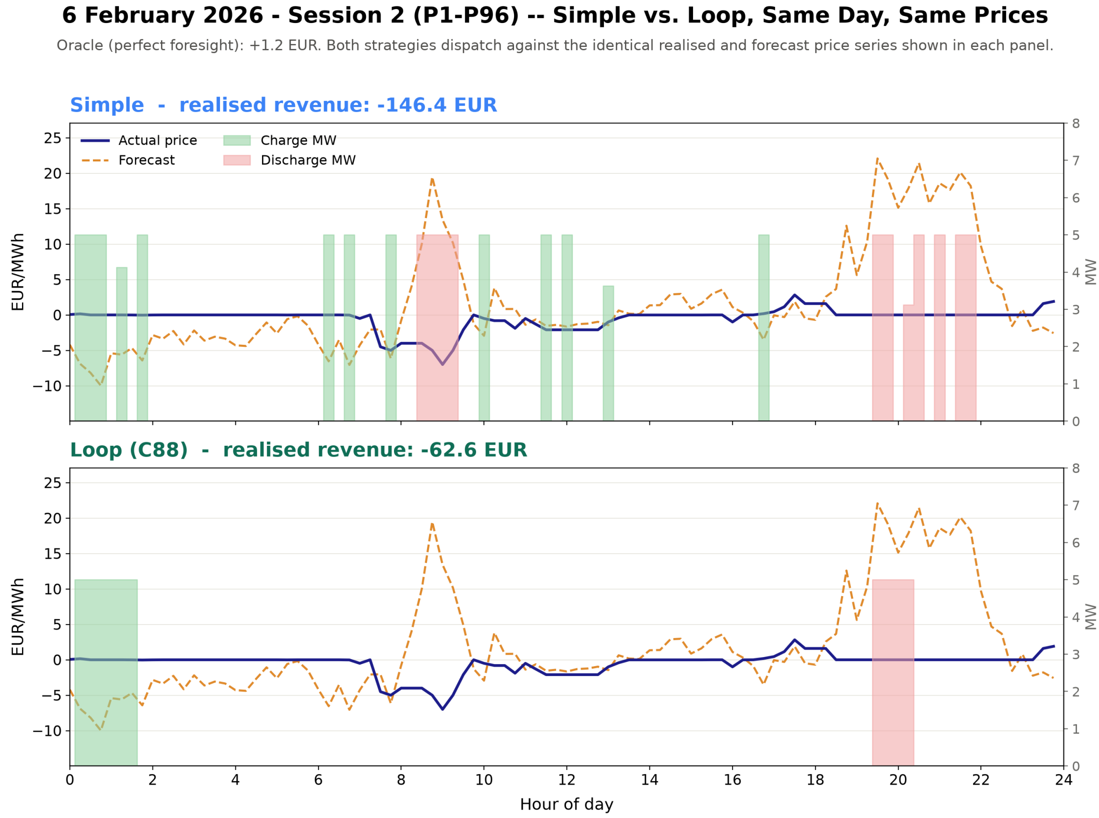

Figure 5.26: 6 February 2026, Session 2: Simple (top) against Loop (bottom) on the identical realised and forecast price series. The realised price (blue) is flat at approximately 0 €/MWh for the entire session, with a single brief dip to −6 €/MWh around 08:00–09:00. The LEAR forecast (orange dashed), by contrast, predicts a sharp, isolated spike to 19 €/MWh in that exact window. Simple's LP still fires eleven charge and five discharge blocks chasing noise across the flat session (top panel), while Loop's single declared pairing is timed precisely to the forecast spike and finds nothing there either (bottom panel): direct visual confirmation that the failure is a forecast-shape error shared by both order types, not a timing or direction error unique to Simple's fragmentation. Both models looked in exactly the right place and found nothing there. This is the same 6 February date documented for Session 1 in Section 5.3.7, confirming the regime shift was market-wide rather than session-specific.

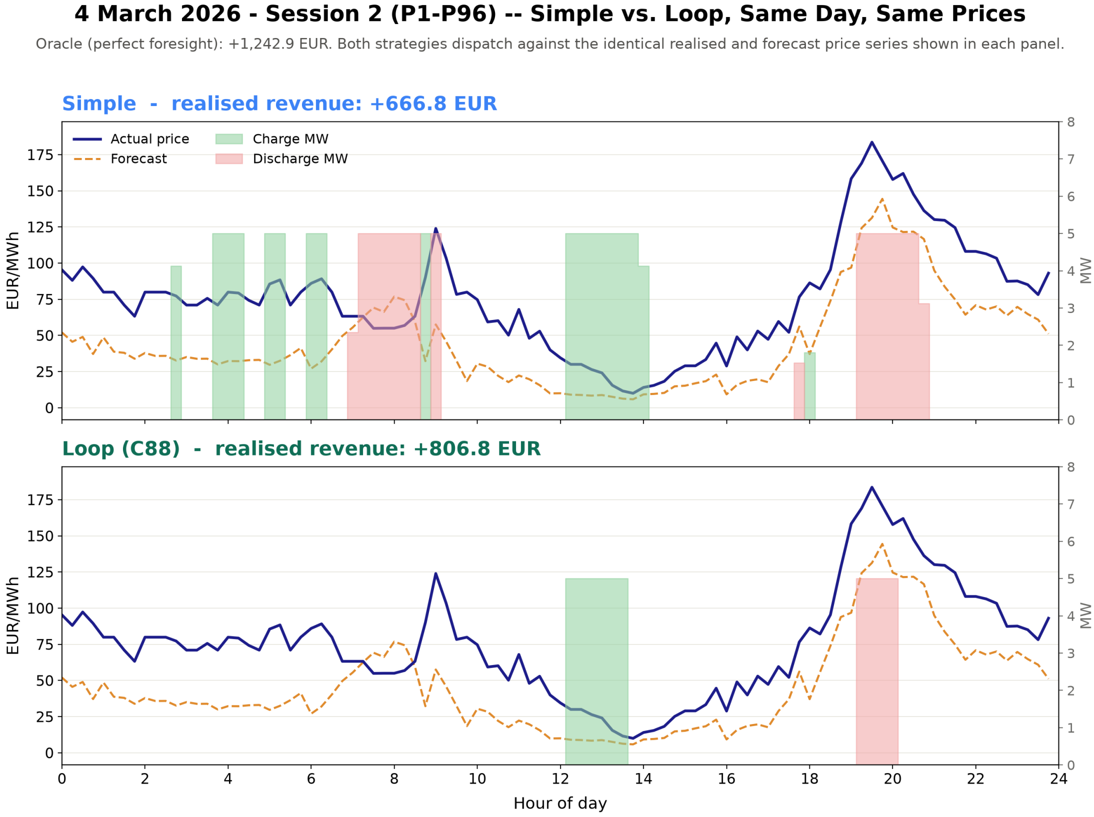

Figure 5.27: 4 March 2026, Session 2: Simple (top) against Loop (bottom) on the identical realised and forecast price series. Loop (806.8 €) outperforms Simple (666.8 €) using a single full-duration cycle: the rule-based search charges through the entire midday solar trough (12:00–14:00) and discharges through the full evening peak (19:00–20:15), rather than fragmenting either side into shorter pulses the way Simple's LP does across the rest of the session (top panel). This is the largest single-day Loop win in the sample, yet both strategies remain well below the 1,242.9 € oracle, since even a full-duration paired bid and Simple's LP schedule each capture only part of the session's available spread.

Why Loop wins here. The underlying dispatch numbers make the mechanism concrete. Simple's LP, across the full session, buys 18.00 MWh of energy for 1,154.8 € (an average price of 64.2 €/MWh paid across nine separate charge blocks) and resells it for 2,001.5 € (111.2 €/MWh average across four discharge blocks), netting 666.8 € after 180.0 € of degradation on the full discharged volume. Loop, by contrast, ignores every period outside its single declared pair: it charges 8.75 MWh over 12:15–14:00 at 21.1 €/MWh (the cheapest contiguous window of that length in the session, deep in the solar trough) and discharges 6.25 MWh over 19:15–20:30 at 168.6 €/MWh (the richest contiguous window following it), for a realised spread of 147.5 €/MWh against a break-even requirement of only 11.8 €/MWh. Because Loop trades roughly a third of the energy Simple does, but at a spread more than double Simple's own average, its realised profit (806.8 €) exceeds Simple's despite touching far less of the session. Simple loses here not because its LP is wrong, but because it is a genuine price-taker across every period, including the seven charge blocks and three of the four discharge blocks priced well inside the trough and peak rather than at their extremes; averaging in those less favourable legs pulls its overall spread down to 47.0 €/MWh, roughly a third of Loop's. Loop wins on days like this one precisely because it is allowed to be selective in a way Simple's full-session optimisation is not: Simple must clear a schedule that nets to positive value across the whole session, while the rule-based Loop search only ever proposes the single best-matched pair it can find and settles for nothing on any other period.

The 23 June 2025 Session 3 near-oracle day is already shown in Figure 5.21 above.

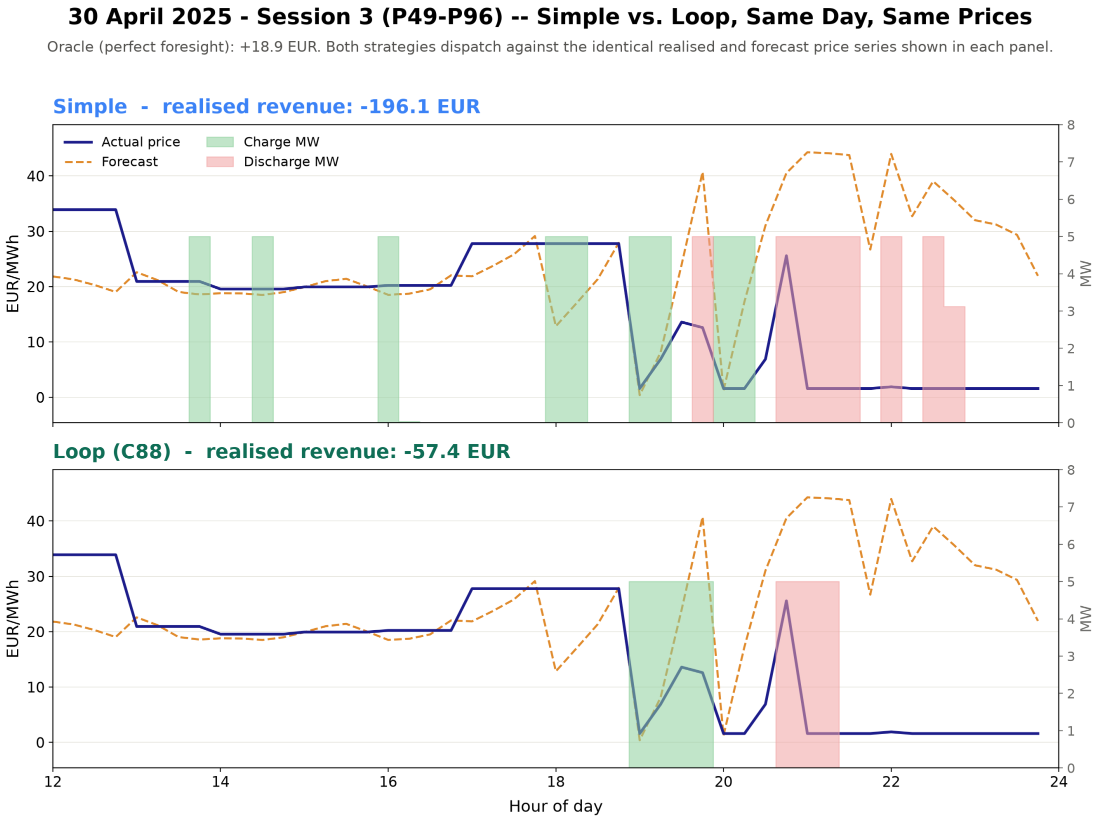

Figure 5.28: 30 April 2025, Session 3: Simple (top) against Loop (bottom) on the identical realised and forecast price series. The realised price (blue) contains a sharp, narrow spike to 26 €/MWh at approximately 20:45. Simple's discharge blocks straddle the spike across several shorter windows from 19:00 to 23:30 without any single block covering it precisely (top panel), while Loop's one declared pairing discharges 20:45–21:30, landing almost exactly on the spike (bottom panel), yet even this well-timed pairing still loses money (−57.4 €), a smaller loss than Simple's (−196.1 €) but a loss nonetheless, because the charge leg bought energy across 19:00–20:00 at prices that, once degradation is included, the single evening spike was not quite rich enough to recover. The LEAR forecast (orange dashed) anticipates generally elevated evening prices but places its own peak around 21:00–22:00, offset from the real spike by roughly 30–45 minutes; this is a timing-resolution failure distinct from the Session 1 Type III example (Figure 5.22), and one that a tighter pairing window would not have fully solved for Loop either, since the charge side was priced against the same imperfect forecast.

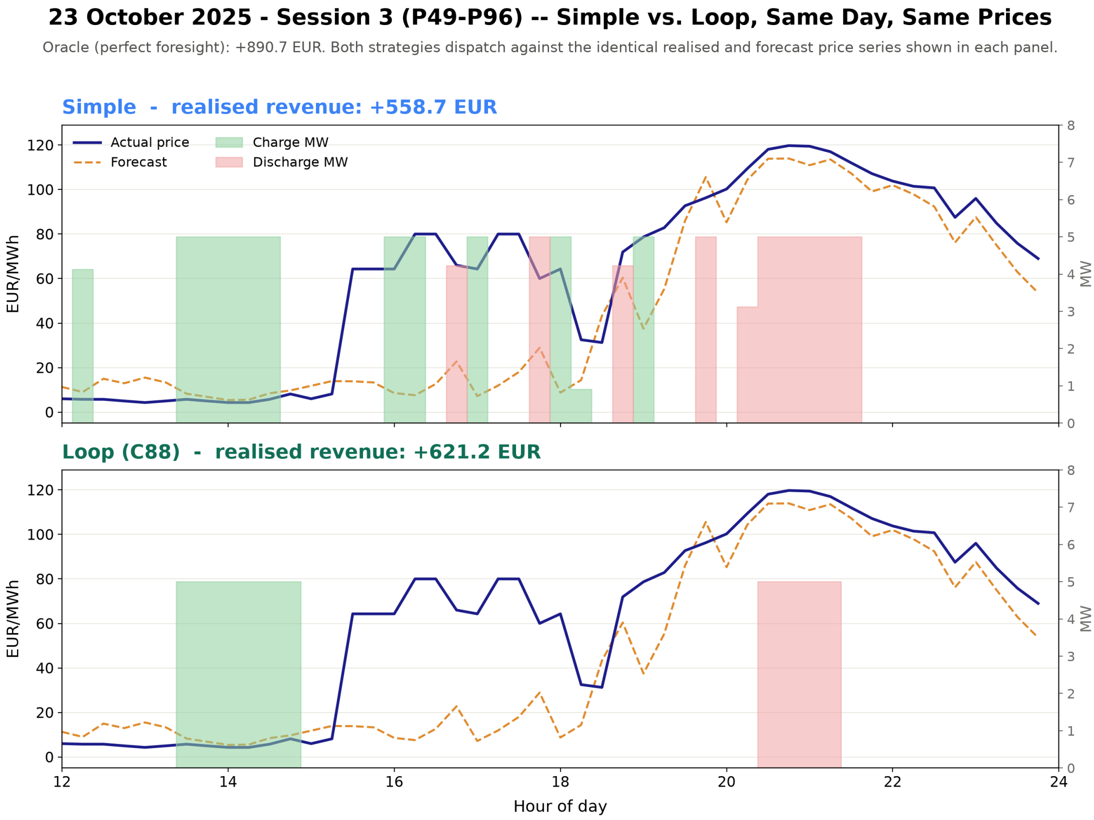

Figure 5.29: 23 October 2025, Session 3 (P49–P96): Simple (top) against Loop (bottom) on the identical realised and forecast price series. Loop (621.2 €) outperforms Simple (558.7 €) using a single full-duration cycle: a continuous charge block through the early-afternoon trough (13:00–15:15) paired with a continuous discharge block through the evening peak (20:30–21:45), against Simple's six charge and five discharge blocks spread across the same session (top panel). Both strategies remain well below the 890.7 € oracle: a genuine, representative Session 3 Loop win, not an outlier.

A second example, cross-session. The same mechanism repeats in Session 3. Simple buys 11.65 MWh across six charge blocks for 484.3 € (41.6 €/MWh average) and resells it across five discharge blocks for 1,159.5 € (99.5 €/MWh average), a 58.0 €/MWh average spread, netting 558.7 € after 116.5 € of degradation. Loop instead charges 8.75 MWh over 01:30–03:15 at 5.6 €/MWh, the overnight trough, well before Simple's own first charge block, and discharges 6.25 MWh over 08:30–09:45 at 117.3 €/MWh, the morning demand ramp; the realised 111.6 €/MWh spread is nearly double Simple's session-wide average, and again more than enough to clear the 11.8 €/MWh break-even margin with room to spare. The pattern across both examples is the same: Loop is not finding revenue Simple misses, since Simple's LP already dispatches into every profitable window in the session, including the ones Loop selects. Loop wins by declining to average its return down over the rest of the session the way Simple's full dispatch necessarily does. This also explains why the technique fails on spike days (Section 5.3.7): Loop needs a window, not an instant, to be worth declaring, so a single sharp price excursion lasting one or two periods is diluted by the average over whatever window contains it, while Simple can bid that exact instant with a single quarter-hourly order and lose nothing to averaging.

The two-cycle case in Table 5.12, 30 October 2025 (Session 1), also shows why a second cycle is harder to realise than the first. Its second, morning leg (04:15–05:00 charge, 09:00–09:30 discharge) was accepted on the forecast, which anticipated a 23.2 €/MWh spread comfortably above the clearing threshold. At actual prices the charge window came in far higher than forecast (2.8 €/MWh forecast versus 89.9 €/MWh realised), and because that leg's charge window (3.75 MWh) is larger than its efficiency-sized discharge window (2.5 MWh), the realised 104 €/MWh discharge price was not enough to cover the cost of the larger charge volume: 337.1 € spent charging against 259.7 € earned discharging, a genuine −102.4 € realised loss on the second leg alone, which is subtracted from the first cycle's own 475.0 € to leave the day's total loop revenue at 372.6 €. The tied charge-and-discharge structure means both sides of this second leg cleared together, exactly as designed; it simply cleared at a loss, since the C88 condition constrains the two legs to execute jointly but does not constrain the price either one clears at. A second cycle is a second independent forecast bet, graded on its own realised outcome rather than the first cycle's: when it is wrong, it actively costs money, not nothing, which is part of why only 13.2 % of Loop's wins in this sample use a genuine second cycle (Section 5.3.10) rather than more, and why two-cycle orders are disproportionately represented among Loop's own loss days.

These days confirm, with the same level of dispatch-level detail across all three sessions, that the fragmentation behaviour driving Simple's own dispatch advantage is a general property of the market, not specific to any one session: both Session 2 examples (Figures 5.25 and 5.27) show the Simple LP splitting its charge allocation into several small windows even when a single trough or ramp dominates the session, while the rule-based Loop instead claims the whole trough and the whole peak in one continuous block each. The Session 2 failure day (Figure 5.26) shows the 6 February 2026 event striking at the same forecast-spike location as in Session 1, confirming it as a single market-wide episode rather than an artifact of one session's forecast window.

5.3.11 | Loop as a Risk-Reduction Instrument

Beyond its direct revenue advantage on selected days, Loop C88's tied charge-and-discharge structure means both legs clear together or neither does; this removes the risk of an unmatched single-sided execution, but it does not by itself guarantee that a cleared pair earns a positive outcome. The C88 condition $\lambda_{h_2} - \lambda_{h_1} \geq c_{\text{deg}}/\eta$ compares the realised prices against the operator's own declared price limits per leg; a pairing built on a forecast that turns out wrong can clear both legs exactly as designed and still fail to cover degradation, in the same way a badly-priced Simple order can. Loop's actual behaviour on Session 1's 18 Simple loss days shows this directly.

On 3 of the 18 Session 1 Simple loss days (30 April 2025, 3 February 2026, 18 February 2026), Loop achieves a positive realised revenue that beats Simple outright. On the remaining 15 days, Loop's own realised revenue is also negative: on 12 of these its loss is smaller than Simple's, so Loop still reduces the size of the loss without eliminating it (for example, 6 February 2026: Loop −37.5 € against Simple's −92.1 €), but on the other 3 (30 December 2025, 1 February 2026, 11 February 2026) Loop's loss is as large as or larger than Simple's – most notably 30 December 2025, where Loop (−62.1 €) loses substantially more than Simple (−6.1 €) despite Simple itself barely losing money that day. In no case across the 18 days does Loop's realised outcome exceed the oracle.

Combining the 3 outright wins with the 12 days where Loop's loss is smaller than Simple's, Loop fully avoids or reduces the Simple loss on 15 of 18 days (83.3 %), but is neutral or actively worse on the remaining 3. This is a materially weaker result than a guaranteed floor: Loop's forecast-driven scheduling reduces exposure to the same regime-shift risk that drives Simple's losses on most days, but it shares that risk rather than being immune to it, since both instruments are built on the same underlying LEAR forecast, and on a minority of days Loop's own forecast is wrong in a way Simple's is not.

Table 5.13 lists all 18 Session 1 Simple loss days with Loop's actual realised outcome alongside. On most of these days the oracle is also very low (below 40 € in most cases, with several at or near 0 €), confirming the market offered little or no genuine spread even under perfect foresight; Loop's narrower, forecast-filtered exposure to these days still leaves it a real loss on most of them, just usually a smaller one than Simple's.

Table 5.13: Session 1's 18 Simple loss days, with Loop's actual realised outcome. Green rows indicate days where Loop achieved a positive realised revenue; yellow rows indicate days where Loop still lost money but outperformed Simple; red rows indicate days where Loop performed worse than Simple.

Date	Simple (€)	Loop (€)	Oracle (€)
19 Apr 2025	−75.2	−22.7	21.6
30 Apr 2025	−40.8	+5.8	140.1
30 Dec 2025	−6.1	−62.1	38.0
31 Jan 2026	−93.1	−84.5	36.1
01 Feb 2026	−71.7	−72.0	0.0
03 Feb 2026	−2.8	+10.7	43.9
05 Feb 2026	−35.9	−14.3	0.0
06 Feb 2026	−92.1	−37.5	0.0
07 Feb 2026	−43.0	−9.5	0.5
08 Feb 2026	−59.1	−6.3	0.1
09 Feb 2026	−56.8	−46.8	11.8
11 Feb 2026	−22.1	−22.7	0.0
13 Feb 2026	−32.1	−21.1	13.4
14 Feb 2026	−27.6	−25.0	24.7
15 Feb 2026	−36.5	−31.5	0.0
16 Feb 2026	−46.4	−30.2	11.6
18 Feb 2026	−23.5	+6.7	21.9
01 Mar 2026	−83.6	−71.9	24.3

The oracle confirms the market offered minimal or no spread on most of these days, so neither order type could have earned much even with perfect foresight. Loop's forecast-driven submission filter consisted of only proposing a pairing when the LEAR forecast shows a positive expected margin. This earns a positive outcome outright on 3 of the 18 days and reduces the size of the loss on a further 12, but it is not a guaranteed floor: on 3 days (30 December 2025, 1 February 2026, 11 February 2026) Loop's own loss is larger than Simple's, and on no day across the 18 does Loop return exactly zero. On the same regime-shift days that break Simple's forecast, Loop's own forecast is typically wrong too. So a genuinely bad session still costs the battery money under Loop, although usually, but not always, less than under Simple.

5.3.12 | Adding a lower bound floor to Loop

Section 4.8.2 established that C88's clearing condition is a declared-price-limit test rather than a break-even test: a leg clears once realised prices beat the operator's own declared bid, independent of whether the resulting margin covers the battery's degradation cost. This is not the only clearing philosophy available in this market. OMIE's own Minimum Income Condition (MIC) already establishes, for conventional generation, that a bid can be withdrawn outright rather than settled at a loss when its realised income falls short of a declared minimum, a precedented reject-on-loss mechanism this thesis proposes extending to Loop. Two results outside this thesis support the idea. First, that price uncertainty threatens a paired storage order's profitability is independently confirmed in the wider literature: a 2023 study of loop-block bidding for storage participants [52] motivates its own model by citing evidence that imperfect price forecasts damage storage profitability more severely than they damage conventional generators, since a storage asset is exposed on both the buy and the sell side at once. Even though Loop's losses are less pronounced than Simple's in our model, it is true that it still faces them, meaning it can still be improved. Second, that MIC's income-condition logic can be modelled in terms compatible with a block order's own clearing surplus is not a novel mathematical claim either: Savelli et al, working on exact linearisations of OMIE's real MIC constraint, note as ongoing work that the income condition could be modelled in a similar way to a block order's surplus [53]. An MIC-style floor for Loop follows directly from both results: the risk it would address is real and documented independently of this backtest, and attaching an income-style condition to a block-shaped instrument already has precedent in the literature.

Real, backtested Loop clears against declared price limits and has no such floor: the 28 negative session-days documented below are exactly what today's Loop actually settles for when a submitted pairing clears but fails to cover degradation. What follows is a bound, not an implementation: *if* Loop's own clearing rule were replaced by an MIC-style reject-on-loss test, rejecting a session's submitted legs outright, at zero revenue, whenever they would otherwise net a realised loss, a session that currently loses money under real Loop would instead settle at zero, and every currently-positive session would be left unchanged. Table 5.14 recomputes each session's Loop total on this modified-clearing-rule basis, using the same 387-day combined sample and burn-in convention used throughout this chapter. It bounds what adding MIC-style clearing to Loop would be worth, not what Loop as backtested elsewhere in this thesis already delivers.

Table 5.14: Bound on the value of adding an MIC-style reject-on-loss clearing rule to Loop, a hypothetical modification of the backtested order type: rule-based Loop's own realised revenue is rejected outright (zero revenue) on every session in which it currently loses money under its real, price-limit clearing rule, and left unchanged otherwise.

Session	Loop total (€)	MIC-style bound (€)	Recovered (€)	Share of Loop total	Negative days
Session 1	157,841	158,479	638	0.4 %	19
Session 2	143,899	146,362	2,462	1.7 %	43
Session 3	117,706	122,487	4,781	4.1 %	56
Combined (387 days)	416,969	419,709	2,740	0.7 %	28

Adding this floor to Loop recovers 2,740 € across the combined sample, a 0.7 % increase over what real Loop already earns. The gain is modest, and for a specific reason: only 28 of 387 combined session-days realise a loss under Loop's existing price-limit clearing to begin with, so there is simply less loss left for an added floor to remove. Loop's own forecast-driven submission filter, which only proposes a pairing when the expected margin across scenarios looks favourable, already does much of this work before settlement happens.

The same underlying idea, applied to Simple's own realised schedule rather than to Loop's, gives the first concrete number for the more ambitious instrument this section has been building toward: a genuinely tied, multi-leg order. Simple's LP dispatch fragments into a median of 10 charge/discharge blocks per session-day, far more than the one or two legs Loop's structure can hold, and has no submission filter of any kind, so it loses money far more often than Loop does. A tied order that could reject one losing pair on a given day while keeping a profitable pair on the same day would require designing and backtesting a genuine multi-leg candidate search, which this thesis has not done. What can be computed exactly, without building that search, is the coarsest possible case of the same instrument: treat Simple's own realised schedule as a single all-or-nothing family, exactly the granularity at which OMIE's real MIC evaluates a generator's entire day, and reject it outright on every day it currently loses money. Table 5.15 gives the result.

Table 5.15: Bound on the reject-on-loss component of a tied, multi-leg order's coarsest possible case: Simple's own realised schedule, treated as a single all-or-nothing family and rejected outright (zero revenue) on every day it currently loses money, left unchanged otherwise.

Session	Simple total (€)	Tied bound (€)	Recovered (€)	Share of Simple total
Session 1	232,551	233,400	849	0.4 %
Session 2	210,626	213,402	2,776	1.3 %
Session 3	165,316	169,664	4,347	2.6 %
Combined (387 days)	604,020	611,307	7,287	1.2 %

Applied this way, the reject-on-loss rule recovers 7,287 € (1.2 %) across the combined sample, nearly twice the relative gain the same rule delivers when applied to Loop's own schedule (Table 5.14). Read together, the two tables make one point from two angles: a reject-on-loss rule recovers a real but modest amount of revenue regardless of which schedule it is applied to, and the effect is largest precisely where the underlying schedule already loses money most often, which in this backtest is Simple's full-session dispatch, not Loop's already-selective, mostly-positive window search. Both figures should be read the same way: they hold today's scheduling logic fixed and credit it only with the right to reject a realised loss outright, at the coarse, whole-session granularity of a single all-or-nothing family rather than a single leg.

This is precisely why the 1.2 % figure is a lower bound on the tied order's value and not an estimate of it. A whole-session test can only accept or reject an entire day at once, never one losing pair while keeping a profitable pair on the same day, so it necessarily misses the larger and structurally more important source of a tied order's value: representing Simple's full multi-leg schedule at all, rather than one all-or-nothing block. Capturing that would require designing the multi-leg candidate search itself and re-evaluating which schedules a per-pair, realised-margin rule would accept as a tied family, which is left for future work (Section 5.3.14). What Table 5.15 already establishes, without that further work, is that the underlying mechanism has a real, computed, non-zero floor: 7,287 € is not a projection or a claim, it is what the coarsest version of the proposed instrument would already have recovered from data already in this backtest.

5.3.13 | Day-Ahead Market as a Validation Cross-Check

The same LEAR + LP + Monte Carlo pipeline was first run on the OMIE day-ahead auction over the same window, using the identical battery model and order-type framework, to confirm that the machinery behaved sensibly on the better-documented day-ahead session before it was applied to the intraday auctions. That run served its purpose as a pipeline check and is not carried further as a revenue comparison: because each market resets the battery to $S_{\min}$ independently, a day-ahead-plus-IDA revenue figure would credit the battery with a fresh usable-energy budget in every market rather than the single physical budget a real unit has for the day, so any such comparison would overstate rather than measure the value of re-optimising across sessions (Section 5.3.14). The day-ahead standalone figures are therefore not reported here; the run is retained in the analysis code.

5.3.14 | Limitations

Several features of the backtest design limit the scope of the conclusions.

The independent-session architecture resets the SoC to $S_{\min}$ at the start of each session and enforces terminal neutrality within it. A real BESS operator would optimise across the day-ahead market and all IDA sessions jointly, exploiting a single evolving SoC. Charging cheaply in Session 1 and discharging profitably in Session 2 generates value the independent model cannot capture. The present design isolates stand-alone session value precisely but understates the total revenue attainable under a coordinated multi-session strategy.

Loop benchmark consistency. is limited by the number of loops the model can bid per session: the battery supports two full charge/discharge cycles per day. A session with two well-separated ramps (an overnight cycle and a separate midday-to-evening cycle, for example) therefore requires the Loop search to consider a second, non-overlapping loop in the remaining free periods and report the combined revenue when one is found. The rule-based method's two-cycle search enforces real-time state-of-charge feasibility by construction: a second cycle is confined either entirely before the first cycle's charge window or entirely after its discharge window, so the battery is never asked to hold both cycles' energy simultaneously. The structurally analogous two-cycle logic in the day-ahead Loop optimiser (`3_order_optimizer.py`) does not enforce this temporal ordering, only excluding the exact periods the first loop occupies when searching for a second; checking the running state of charge this implies shows it can schedule the battery as holding more energy than its 9 MWh usable capacity at some point in the session. This is one further reason the day-ahead cross-check (Section 5.3.13) is used only to validate the pipeline and is not reported as a revenue comparison here.

The February 2026 renewable surplus episode illustrates a structural limitation of the LEAR architecture: because the training window expands rather than discards older observations, a sustained, novel market regime is heavily outweighed by the larger volume of historical “normal” days already in the training set, and the 14-day rolling window used for the dispersion estimate $\hat{\sigma}_t$ is too short-sighted on its own to compensate. Neither mechanism adapts quickly enough to a regime this different from the training history. Figure 5.23 (6 February 2026, Session 1) is a representative instance: the LEAR forecast still anticipates a normal-spread day while the actual IDA price sits near zero throughout, so the model dispatches into a market with no real spread and loses 92 € in degradation costs it cannot recover. This is the Type I (phantom-spread) failure mode identified in Section 5.3.7, and it accounts for 13 of the 18 Session 1 loss days. A model that incorporates regime-change detection or a longer adaptive window would reduce these losses, at the cost of slower adaptation during normal periods. A related but distinct limitation is illustrated by Figure 5.22 (30 April 2025): here the forecast correctly identifies that a spread exists but locates the charge window incorrectly, again eroding the realised margin below the degradation threshold (Type III failure). By contrast, the Type II failure mode that dominates the day-ahead backtest’s largest single-day losses — the model missing an extreme price spike entirely and submitting no bid, discussed in Section 5.3.7 — does not appear in the post-burn-in IDA sample. This should be read as a property of the shorter intraday forecast horizon rather than evidence that the intraday model is generally more robust: the absence of Type II failures in this specific 387-day window does not guarantee they cannot occur in a longer or differently-timed sample.

The degradation cost is fixed at 10 €/MWh discharged throughout. This is a limitation which specifically concerns Simple’s fragmented dispatch pattern documented throughout Section 5.3 (e.g. Figure 5.20, where a single trough is split into three separate charge pulses). At the power-electronics level this is physically achievable: a battery inverter can reverse direction between one quarter-hour period and the next with no meaningful constraint, so fragmenting a schedule into many short windows rather than one continuous block is not infeasible the way the Loop SoC violation identified in Section 5.3.2 was. The simplification lies in the degradation cost model itself: it charges a flat 10 €/MWh on total energy discharged, regardless of how that energy was accumulated, so one continuous four-hour discharge costs exactly the same as eight scattered fifteen-minute discharges totalling the same energy. Real battery degradation is more nuanced than a pure throughput charge – cycle-counting effects (many small partial cycles are not always equivalent to fewer deep cycles, depending on cell chemistry) and switching-related losses at each charge/discharge transition are both absent from the present model. Since this affects Simple’s fragmented schedules more than Loop’s one- or two-block structure, correcting it would be expected to narrow, not widen, the gap between the two order types; quantifying the effect would require a degradation cost model sensitive to cycling pattern, not only to total throughput, which is left for future work.

Finally, the backtest treats IDA prices as exogenous. A 5 MW price-taker approximation is defensible at the scale of batteries present now and in the following years in Spain’s electricity wholesale market, thus the model does not account for potential price impact from the battery’s own orders or for strategic interactions with other storage units bidding in the same session. However, more batteries entering the generation mix will affect the prices both in the day-ahead and the intraday market. This would need to be revised in the following years.

6 | Conclusions

OMIE's intraday auctions currently offer two instruments – simple orders and simple blocks – a product set that, like every NEMO's, reflects the operating requirements of the technologies that have historically dominated its market. Section 5.2 showed empirically that OMIE's existing block order is used almost exclusively by combined-cycle gas turbines guaranteeing minimum run windows and by pumped hydro coordinating charge-discharge pairs, at sizes one to two orders of magnitude above a typical battery's rated power. The instrument serves the assets it was designed for well; what it was not designed for, because battery deployment in Spain has only recently begun to accelerate, is the round-trip economics of newer storage assets at a fraction of that scale. This thesis set out to establish, from OMIE's own published data rather than from economic theory, what the intraday order-type catalogue does for the assets that use it today and what it fails to do for the asset class arriving next – and then to quantify that failure in euros rather than assert it in principle.

What the empirical record shows

The first half of the contribution is descriptive, and it stands independently of anything the thesis concludes about batteries. Across the 365-day sample, block orders account for just 2.25 % of IDA submissions by count yet 9.7 % of cleared megawatt-hours, and the instrument is used by a narrow, structurally coherent set of participants: CCGTs and pumped hydro together supply 87.4 % of all block submissions, while nuclear, wind, and solar – which together submit over 1.8 million IDA offers across the sample – reach for a block on fewer than 0.1 % of them. That contrast is the empirical heart of the descriptive finding. Technologies with a genuine physical indivisibility – a thermal minimum-run window, a pumped-hydro filling cycle – adopt the block; technologies that can ramp freely and express their constraints through ordinary price-quantity curves ignore it entirely, despite being fully active in the same sessions. The block order is not a general-purpose complex instrument that happens to be under-used; it is a targeted response to a specific class of operating constraint, and the offer data show participants treating it exactly that way.

A second empirical regularity gives that picture a mechanism. How far a block is priced from the market it actually faced – the average realised `marginalpibc` price over the block's own delivery span – predicts clearance almost deterministically, and the relationship is monotonic across every technology/direction pair with a meaningful sample: pumped hydro's sell blocks, priced on average below the market they saw, clear three times in four, while coal's sell blocks, standing roughly 100 €/MWh above their realised market, clear under one time in a hundred. This is not miscalibration. A block priced far above the market is a standing option at a floor the unit does not expect the market to reach, a deliberate hedge on a minimum-generation constraint rather than an attempt to trade volume, and its low clearance rate is the intended outcome, not a failure of the bidder. Read together, these two findings reframe what a low aggregate usage or clearance rate means for an intraday product: they are not evidence that an instrument is redundant, but the signature of an instrument doing a narrow job for the participants who need it. That reframing matters directly for the design question the thesis ultimately asks, because it dissolves the intuition that a rarely-used or rarely-cleared instrument is a candidate for removal.

What the backtest shows

The second half of the contribution is quantitative. The backtest in Section 5.3 put a number on the mismatch between the existing catalogue and a battery's operating profile. Across 387 operative days and three intraday sessions, a battery bidding the best available order type each session captured 81.4 % of the perfect-foresight welfare benchmark, leaving an average gap of 18.6 %. The welfare-gap decomposition is what makes that figure interpretable rather than merely reported: almost all of the gap is ordinary forecast error – the price a battery pays for committing before it knows the session's clearing price – and not order-type constraint. The failure-mode taxonomy sharpens this further. The single largest contributor to Session 1's shortfall is a concentrated cluster of phantom-spread days in early February 2026, when a sustained renewable surplus pushed IDA prices to near zero and the LEAR forecaster, drawing on an expanding training history dominated by normal-spread days, dispatched into a market that offered nothing to arbitrage. That the largest losses trace to a novel price regime the forecaster had not yet seen, rather than to any structural inadequacy of the order type, is itself a finding: it locates the residual gap in forecast quality and regime detection, which a better price model can attack, rather than in the bidding instrument, which it cannot.

A paired charge-discharge order modelled on EPEX SPOT's Loop Block Order (C88), scheduled with a rule-based heuristic that targets the battery's full per-cycle duration rather than the single sharpest average price (Section 5.3.2), wins outright on 5.0 %, 7.6 %, and 4.6 % of Session 1, 2, and 3 days respectively, and more than doubles Loop's realised revenue relative to a naive expected-value ranking in every session. The lesson generalises beyond this thesis: for a paired-block instrument, the scheduling rule is not an implementation detail but a first-order determinant of whether the instrument looks viable at all, and an expected-value ranking is actively the wrong objective for it.

Even under the better scheduling method, Loop is not a revenue-maximising alternative to Simple bidding: it beats Simple's period-by-period flexibility on fewer than one day in ten. That negative result is the most informative single outcome of the backtest, because of *why* it holds. Simple's LP dispatch fragments into a median of 10 separate charge and discharge blocks per session-day, and 98.7 % of session-days require more than the two legs a single Loop pairing can express. A battery's price curve, at quarter-hourly resolution, is essentially never a single clean ramp; it is a sequence of many small opportunities, and an instrument that can capture only one of them, however well chosen, is structurally outmatched by one that can capture all of them. This is not a contingent fact about Spanish prices or about the particular battery modelled – it follows directly from pairing an instrument shaped around the single-cycle economics of thermal plant and pumped hydro with an asset that has neither constraint and can reverse direction every fifteen minutes. The very rigidity that makes the block order valuable to a CCGT – its all-or-nothing commitment across a contiguous span – is precisely what makes it ill-suited to an asset whose value lives in the freedom to fragment.

Loop nonetheless earns a place through a narrower channel that the expected-revenue comparison alone hides, though the exact nature of that channel needs to be stated precisely. Loop's tied charge-and-discharge structure guarantees only that both legs clear together or neither does; it does not guarantee that a cleared pair earns a positive outcome, since the C88 condition compares realised prices against the operator's own declared limits per leg, not against a break-even threshold. On Session 1's 18 Simple loss days, Loop wins outright on 3 and reduces, but does not eliminate, the loss on a further 12 by realising a smaller loss of its own; on the remaining 3, Loop's own loss is as large as or larger than Simple's. A Simple order has no equivalent mechanism: once both of its legs clear at their marginal prices, the battery pays degradation whether or not the round trip was profitable, whereas Loop's forecast-driven submission filter – proposing a pairing only when the LEAR forecast shows a positive expected margin – reduces exposure to the same regime-shift risk that drives Simple's losses without eliminating it, since both instruments share the same underlying forecast. Loop's genuine contribution in this backtest is therefore partial downside mitigation, with a modest revenue contribution on a small minority of days – not a routine alternative to Simple, and not a guaranteed floor, but a conditional, imperfect hedge whose value shows up mainly in the size of the loss avoided rather than its complete elimination.

The day-level evidence (Section 5.3.8) is specific enough to say when an operator should actually prefer Loop over Simple, rather than treating the choice as a coin flip decided after the fact. Three conditions recur across the winning days examined in detail. First, Loop wins when the session's price structure reduces to one dominant, sustained ramp – a single trough deep enough and a single peak rich enough that the average price across a multi-period window still clears a wide margin – rather than several smaller opportunities spread across the session; concretely, this means low-volatility, solar-driven shoulder-season and winter days (67 of 68 wins fall in February–May or October–January), not the sharp, isolated spikes that a spread-chasing intuition would expect to favour a block order, since Loop's window-average settlement dilutes a spike the way it never dilutes a sustained ramp – visible directly in the 4 March 2026 and 23 October 2025 case studies (Figures 5.27 and 5.29), where Loop claims the full trough and the full peak in one continuous block each while Simple fragments the same ramp into several smaller legs. Second, Loop is the more cautious choice specifically when forecast confidence is low: because it only ever proposes a pairing when the LEAR forecast shows a positive expected margin, an operator facing a session where the price regime looks uncertain – shortly after a structural shift like October 2025's move to quarter-hourly scheduling, or during a period the forecaster has not seen the likes of before, such as the February 2026 renewable-surplus episode – gives up Simple's higher expected revenue for a submission rule that reduced the size of the loss on 15 of Session 1's 18 loss days, though it eliminated the loss entirely on only 3 of them and was neutral or worse on the remaining 3; this is a partial hedge, not a guarantee. Third, the decision of which order type to submit is a per-session forecast-confidence judgement rather than a fixed rule: on a session where the price regime looks uncertain, the operator accepts Loop's lower expected revenue for its partial loss mitigation; on a forecast-confident, single-ramp session, Simple's full flexibility remains the higher-expected-value choice on 90.4 % of days in this backtest, so the case for choosing Loop over Simple is the exception, not the default.

Reading the two halves together

The descriptive and quantitative findings converge on a single reading of OMIE's intraday market. The order-type catalogue is not arbitrary and it is not neglectful: every instrument in it maps onto an identifiable operating constraint of an incumbent technology, and the participants who face those constraints use the instruments exactly as their design predicts. The catalogue's limitation is not internal inconsistency but coverage. It was assembled around the asset classes that dominated Spanish intraday rebalancing for two decades, but it will need to keep adjusting to the necessities of its participants. For this, OMIE will need to consider instruments shaped around constraints that define more recent technologies. In this case the analysis has focused on the newer storage assets, which have the requirement that a charge and its matching discharge succeed or fail together. The block order comes closest, but it encodes that coupling through a rigid, large, single-direction commitment built for a different physical problem, and the backtest shows a battery gains nothing operational from it while inheriting its clearing risk. This is why the gap is best understood as a question of catalogue coverage rather than product quality, and why the appropriate response is addition rather than reform.

That framing also disciplines how strong a claim the evidence can support. A battery captures four-fifths of the available welfare under the existing Simple order, so the missing instrument is not the difference between a functioning and a broken market – it is a targeted improvement at the margin, most valuable in exactly the conditions (thin or adverse spreads, sustained novel regimes) where an unconditional Simple order exposes the operator to avoidable losses. The case for acting on it rests not on the size of the average gap, which forecast quality dominates, but on the existence of a recurring, identifiable market condition that no current instrument lets a battery express – the same standard by which OMIE already justifies offering a block order used on only 2.25 % of submissions. Judged by that standard, and not by a raw win-rate, the evidence for a storage-oriented instrument is affirmative.

A direction for product design

Where the evidence points, then, is neither to defend the block order for batteries nor exclusively leave Simple bidding, but to a narrow generalisation suggested directly by the fragmentation finding. Two paths follow from it, and they ask very different things of OMIE. The first is the possibility of adopting an instrument that already works elsewhere in Europe, optionally sharpened with a rule OMIE already runs for a different purpose. The second is a genuinely new clearing mechanism that neither OMIE nor any NEMO studied in this thesis currently offers.

The simpler step: bring Loop to OMIE, and give it a floor. Loop (C88) is not part of OMIE's current product set. It has been evaluated throughout this thesis as a policy counterfactual: what a battery could earn if an equivalent order existed at OMIE. The first and least demanding recommendation is simply that OMIE adopt it, since the backtest in Chapter 6 already shows its value as a downside hedge even in its plain form, where a leg clears once realised prices beat the operator's declared bid, with no further test of whether the resulting margin actually covers cost.

That plain form, however, could be sharpened further, and cheaply, once Loop exists in OMIE at all. OMIE's own Minimum Income Condition already establishes, for conventional generation, that a bid can be withdrawn outright rather than settled at a loss when its realised income falls short of a declared minimum. Attaching this same logic to a newly introduced Loop order, so that a pair clears only if its realised margin covers the battery's degradation cost rather than merely beating a declared price limit, is not a new market design so much as reusing a rule OMIE already runs elsewhere for a new instrument. Applying it to Loop's real, backtested clearing logic recovers 2,740 € (0.7 %) on top of what an adopted Loop would already earn. The gain is modest, and for a specific reason: Loop's own forecast-driven filter already avoids a loss on most days by simply not submitting when the expected margin looks unfavourable, so there are fewer losing sessions left for a reject-on-loss rule to catch. However, it would make Loop's argument as a risk-hedging instrument even more sound, since a reject-on-loss floor makes it impossible for Loop to lose money outright. Such a floor however helps most exactly where the underlying schedule already loses money most often, and in this backtest that is Simple's full-session dispatch, not Loop's already-selective window search.

The harder step: a tied, multi-leg Simple order. The more demanding path follows directly from the same fragmentation finding that motivates this whole section: a battery's disadvantage under Loop is that it can pair only one or two windows where Simple freely uses ten. Removing that disadvantage means an instrument that keeps Loop's joint clear-or-reject guarantee without inheriting the block structure Loop carries from its thermal and pumped-hydro origins. Concretely, this would be a tied pair of ordinary Simple offers, priced and timed exactly as Simple orders already are, cleared together only if the realised spread between them covers the battery's degradation cost, submitted as many times per session as the price curve warrants. In one sense this is not a new market design at all, only a joint-clearing flag attached to selected pairs of existing orders. In practice, though, it asks far more of the clearing engine than adopting and extending Loop does: EUPHEMIA and the NEMOs would need to evaluate many simultaneous paired constraints per session, not a single instrument with a fixed two-leg structure, which is a meaningfully heavier computational and implementation lift than introducing an order type that already runs at EPEX today.

The evidence for this second step is suggestive rather than conclusive, and the thesis is deliberate about the boundary. The no-trade guarantee alone, applied to Simple's own realised schedule as a single whole-session test (the same coarse granularity at which OMIE's real MIC evaluates a generator's entire day, rather than leg by leg), recovers 7,287 € across the combined sample, a 1.2 % increase in Simple revenue. This is a real but modest floor, and this thesis can bound it exactly because it requires no new search: it is not the tied instrument itself, only a conservative stand-in for it. A whole-session test can only accept or reject an entire day at once, while a true tied, per-pair test could reject one losing pair and keep a profitable pair on the same day. The fully tied instrument's eventual floor should therefore be read as strictly tighter than the 1.2 % figure above, not merely a variant of it. The larger component, representing substantially more of Simple's ten-block schedule than Loop's two legs can hold, cannot be bounded without building and backtesting the multi-leg search itself.

The 1.2 % and 0.7 % figures are therefore lower bounds on the value of a reject-on-loss rule under each order type, not estimates of it. Formally specifying and evaluating the multi-leg tied order through the same welfare-gap architecture used throughout this thesis is the most direct continuation of this work. The contribution here is to establish, from the fragmentation and downside statistics, that both extensions are worth building, and to fix precisely the evidence each would need to clear: a first step that mostly asks OMIE to adopt something that already exists elsewhere, and a second that would require real engineering investment from the market operator to deliver in full.

Methodological contribution and scope

The methodology assembled to reach these conclusions is a contribution in its own right, independent of the Loop-order result. The welfare-gap decomposition – separating forecast-error loss from order-type-constraint loss against a common perfect-foresight benchmark – is not specific to Loop, to batteries, or to OMIE: it generalises to any comparison between a candidate instrument and an incumbent one for any storage technology, and could be applied unchanged to evaluate the tied-order extension above or any other product ACER or OMIE might weigh in a future rulebook revision. It converts the abstract question "is this order type expressive enough" into a measurable quantity with a decomposable structure, which is what lets a market operator reason about product design against evidence. Alongside it, the unit-level block-order characterisation in Section 5.2 is, to the author's knowledge, the first empirical account of how OMIE's existing complex orders are actually used in the intraday auction, and it stands regardless of anything the thesis concludes about storage.

The scope of these findings is bounded in ways worth stating plainly. The session-independence assumption isolates the value of each auction cleanly but understates what a battery could earn by coordinating state of charge across the day-ahead market and all three IDA sessions jointly: resetting the battery to $S_{\min}$ at the start of every session grants it a fresh 9 MWh of usable cycling capacity in each market rather than the one physical envelope a real unit has for the full day, so the present design cannot show how much of that capacity a coordinated, single-SoC strategy would actually have available to re-deploy across markets on the same day. A sequential model that carries a single evolving state of charge through the day is therefore the most consequential extension of this work. The backtest also treats IDA prices as exogenous, a defensible approximation for a single price-taking 5 MW unit but one whose validity erodes as aggregate storage grows toward Spain's 22.5 GW 2030 target; the equilibrium effects of that build-out on OMIE's own price formation lie outside a price-taking model and will need their own treatment as the fleet scales. And the tied, multi-leg instrument the evidence points toward has not itself been implemented here, by design: the thesis establishes the case for it and the standard it must meet, and leaves its construction as the defined next step.

None of these limitations disturbs the central finding. OMIE's intraday product set was designed around the operating constraints of thermal and pumped-hydro assets, and it shows: the instruments that exist are used exactly as their design would predict, and the one that would best serve a cycling battery is not the block order that exists today but a natural generalisation of it: a tied, multi-leg paired order that trades Loop's rigid two-leg structure for something closer to Simple's own dispatch flexibility, while keeping the net-margin guarantee that gives Loop its genuine value in this backtest. As battery deployment in Spain accelerates toward the scale the PNIEC has aimed for, the case for extending OMIE's product catalogue in this direction will only strengthen. The contribution of this thesis is to show precisely where this catalogue has room to grow, to measure the cost of that shortfall against a perfect-foresight benchmark rather than assert it, and to fix the evidence an instrument for the next-generation technologies would need to be judged against.

References

[1] Market Coupling Steering Committee, "Work Plan of the MCSC for 2026 Work Plan of the MCSC for 2026 2," 2025.

[2] Next Kraftwerke, *Market Coupling: How the European electricity market is linked*, https://www.next-kraftwerke.com/knowledge/market-coupling.

[3] European Commission, "COMMISSION REGULATION (EU) 2015/1222 - establishing a guideline on capacity allocation and congestion management," 2015, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32015R1222.

[4] Nemo Committee, "Prerequisites for becoming a NEMO," 2026, https://www.nemo-committee.eu/assets/files/legal-basis-nemo-prerequisites-paper.pdf.

[5] Nemo Committee, "EUPHEMIA Public Description Single Price Coupling Algorithm," 2025, https://www.nemo-committee.eu/assets/files/euphemia-public-description-2020.pdf.

[6] CACM, *CACM 2024 Annual Report*, https://www.nemo-committee.eu/assets/files/cacm-annual-report-2024.pdf.

[7] Nemo Committee, *NEMOs: Nominated Electricity Market Operators*, https://www.nemo-committee.eu/index.

[8] L. H. Andri Busch, *SDAC 15-minute MTU transition: What's changing for European short-term power traders?* 2025, https://dexterenergy.ai/news/sdac-15-minute-mtu-transition/.

[9] EPEX SPOT, *Trading Products*, https://www.epexspot.com/en/tradingproducts.

[10] OMIE, "OMIE DAY-AHEAD AND INTRADAY ELECTRICITY MARKET RULES Non-binding Translation of the electricity market rules," 2014.

[11] EIRGRID, "Industry Guide to the I-SEM I-SEM Project," 2017, http://www.sem-o.com/isem/Pages/Home.aspx.

[12] ENTSOE, *Transparency Platform*, https://transparency.entsoe.eu/.

[13] L. S. Mareike Moraal, *Understanding the German Nuclear Exit | Heinrich Böll Stiftung | Washington, DC Office - USA, Canada, Global Dialogue*, 2023, https://us.boell.org/en/2023/04/21/understanding-german-nuclear-exit.

[14] EPEX SPOT, "Trading at EPEX SPOT (trading brochure)- February 2025," Tech. Rep., Feb. 2025, https://www.epexspot.com/sites/default/files/2025-02/EPEX%20SPOT%20Trading%20Brochure%202025%20February.pdf.

[15] EPEX SPOT, "EPEX Cross-Border Intraday: Questions & Answers," Tech. Rep., 2018, https://www.epexspot.com/sites/default/files/2019-02/XBID%20Q&A.pdf.

[16] EPEX SPOT, *SIDC: Expansion of 30 and 15 minute products and Revised Local Implementation Project 14 Go-Live*, 2020, https://www.epexspot.com/en/news/sidc-expansion-30-and-15-minute-products-and-revised-local-implementation-project-14-go-live.

[17] EPEX SPOT, *Day-Ahead volumes in Austria and Germany-Luxembourg display stable growth*, 2018, https://www.epexspot.com/en/news/day-ahead-volumes-austria-and-germany-luxembourg-display-stable-growth.

[18] EPEX SPOT, "EPEX - Power Spot Auctions," Tech. Rep., 2023, https://www.epexspot.com/sites/default/files/download_center_files/EPEX%20SPOT_Power%20Spot%20Auctions_Block%20Parameters.pdf.

[19] *Italy Electricity Generation Mix 2025 | Low-Carbon Power Data*, https://lowcarbonpower.org/region/Italy.

[20] D. Hine, "The Italian referendums of 8/9 November 1987," *Electoral Studies*, vol. 7, no. 2, pp. 163–167, Aug. 1988, ISSN: 02613794, 10.1016/0261-3794(88)90029-7, https://www.sciencedirect.com/science/article/abs/pii/0261379488900297.

[21] Euronews Green, *Italy set new solar and wind records last year but is still off track for 2030, expert says*, 2024, https://www.euronews.com/2024/01/23/italy-set-new-solar-and-wind-records-last-year-but-is-still-off-track-for-2030-expert-says.

[22] Giulio Maroncelli & Riccardo Malavolti, *Regulatory change in Italy and how it affects PPAs*, 2024, https://www.dlapiper.com/en-es/insights/publications/2024/04/dal-prezzo-unico-nazionale-pun-al-prezzo-zonale.

[23] APE- Appunti di Energia, "IL MERCATO INFRAGIORNALIERO," 2024.

[24] Gestore Mercati Energetici, "Technical Rule no. 5 rev. 05 MPE (TIDE)," Tech. Rep., 2026, https://www.mercatoelettrico.org/Portals/0/Documents/en-US/20260201DTF05rev5MPE_En.pdf.

[25] N. Fabra and M. Reguant, "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, vol. 104, no. 9, pp. 2872–99, Sep. 2014, ISSN: 0002-8282, 10.1257/AER.104.9.2872.

[26] S.-E. Fleten and T. K. Kristoffersen, "Stochastic programming for optimizing bidding strategies of a Nordic hydropower producer," *European Journal of Operational Research*, vol. 181, no. 2, pp. 916–928, 2007, https://ideas.repec.org/a/eee/ejores/v181y2007i2p916-928.html%20https://ideas.repec.org//a/eee/ejores/v181y2007i2p916-928.html.

[27] M. Reguant, "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets," *The Review of Economic Studies*, vol. 81, no. 4, pp. 1708–1742, Oct. 2014, ISSN: 0034-6527, 10.1093/RESTUD/RDU022, https://dx.doi.org/10.1093/restud/rdu022.

[28] L. Hirth, "The Market Value of Variable Renewables," *SSRN Electronic Journal*, Jul. 2012, 10.2139/SSRN.2110237, https://papers.ssrn.com/abstract=2110237.

[29] OMIE, *Presentaciones CAM*, 2026, https://www.omie.es/es/presentaciones-cam.

[30] Silvestros Vlachopoulos (LCP), *Review of the EASE European Market Monitor On Energy Storage (EMMES)*, https://www.lcp.com/en/media-centre/press-releases/europe-accelerates-renewable-energy-growth-89-gw-of-energy-storage-installed-by-the-end-of-2024-with-a-strong-outlook-for-2030.

[31] European Parliament and Council, *Directive 2019/944/EU*, https://eur-lex.europa.eu/eli/dir/2019/944/oj/eng.

[32] R. Weron, "Electricity price forecasting: A review of the state-of-the-art with a look into the future," *International Journal of Forecasting*, vol. 30, no. 4, pp. 1030–1081, Oct. 2014, ISSN: 0169-2070, 10.1016/J.IJFORECAST.2014.08.008, https://www.sciencedirect.com/science/article/pii/S0169207014001083.

[33] J. Lago, G. Marcjasz, B. De Schutter, and R. Weron, “Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark,” *Applied Energy*, vol. 293, p. 116 983, Jul. 2021, ISSN: 0306-2619, 10.1016/J.APENERGY.2021.116983, https://www.sciencedirect.com/science/article/pii/S0306261921004529.

[34] B. Uniejewski, J. Nowotarski, and R. Weron, “Automated Variable Selection and Shrinkage for Day-Ahead Electricity Price Forecasting,” *Energies 2016, Vol. 9, Page 621*, vol. 9, no. 8, p. 621, Aug. 2016, ISSN: 1996-1073, 10.3390/EN9080621, https://www.mdpi.com/1996-1073/9/8/621/htm%20https://www.mdpi.com/1996-1073/9/8/621.

[35] B. Uniejewski and R. Weron, “Efficient Forecasting of Electricity Spot Prices with Expert and LASSO Models,” *Energies 2018, Vol. 11, Page 2039*, vol. 11, no. 8, p. 2039, Aug. 2018, ISSN: 1996-1073, 10.3390/EN11082039, https://www.mdpi.com/1996-1073/11/8/2039/htm%20https://www.mdpi.com/1996-1073/11/8/2039.

[36] B. Uniejewski, G. Marcjasz, and R. Weron, “Understanding intraday electricity markets: Variable selection and very short-term price forecasting using LASSO,” *International Journal of Forecasting*, vol. 35, no. 4, pp. 1533–1547, Oct. 2019, ISSN: 0169-2070, 10.1016/J.IJFORECAST.2019.02.001, https://www.sciencedirect.com/science/article/abs/pii/S0169207019300123.

[37] R. Sioshansi, P. Denholm, T. Jenkin, and J. Weiss, “Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects,” *Energy Economics*, vol. 31, no. 2, pp. 269–277, Mar. 2009, ISSN: 0140-9883, 10.1016/J.ENECO.2008.10.005, https://www.sciencedirect.com/science/article/abs/pii/S0140988308001631.

[38] J. M. M. González, A. J. Conejo, H. Madsen, P. Pinson, and M. Zugno, “Integrating Renewables in Electricity Markets: Operational Problems,” *Springer*, vol. 205, p. 429, 2014, 10.1007/978-1-4614-9411-9, https://orbit.dtu.dk/en/publications/integrating-renewables-in-electricity-markets-operational-problem/.

[39] Jonathan Touriño Jacobo, *Spain increases energy storage target in NECP to 22.5GW by 2030*, https://www.energy-storage.news/spain-increases-energy-storage-target-in-necp-to-22-5gw-by-2030/.

[40] Energy Storage EUrope, “NATIONAL ENERGY & CLIMATE PLANS OVERVIEW OF ENERGY STORAGE AND FLEXIBILITY TARGETS ANALYSIS,” Tech. Rep., 2025.

[41] APPA Renovables, “2025 INFORME anual de autoconsumo fotovoltaico y almacenamiento,” https://www.informeautoconsumo.es/wp-content/uploads/2026/02/Informe-Autoconsumo-Fotovoltaico-Almacenamiento-2025.pdf.

[42] Javier Revuelta and Ignacio Cobo Berberana, “Revista Energética XXI,” *Revista Energética*, pp. 56–57, https://energetica21.com/static/revistas/104/pdf/abril-2026.pdf.

[43] C. O’Connor, M. Bahloul, S. Prestwich, and A. Visentin, “A Review of Electricity Price Forecasting Models in the Day-Ahead, Intra-Day, and Balancing Markets,” *Energies 2025, Vol. 18, Page 3097*, vol. 18, no. 12, p. 3097, Jun. 2025, ISSN: 1996-1073, 10.3390/EN18123097, https://www.mdpi.com/1996-1073/18/12/3097/htm%20https://www.mdpi.com/1996-1073/18/12/3097.

[44] B. Mirletz et al., "Annual Technology Baseline: The 2024 Electricity Update (NREL)," *NREL*, 2024, https://docs.nlr.gov/docs/fy24osti/89960.pdf.

[45] T. Bowen, I. Chernyakhovskiy, and P. Denholm, "Grid-Scale Battery Storage: Frequently Asked Questions (NREL/TP-6A20-74426)," *NREL*, https://docs.nlr.gov/docs/fy19osti/74426.pdf.

[46] B. Xu, A. Oudalov, A. Ulbig, G. Andersson, and D. S. Kirschen, "Modeling of lithium-ion battery degradation for cell life assessment," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 1131–1140, 2018, ISSN: 19493053, 10.1109/TSG.2016.2578950.

[47] Y. Shi, B. Xu, B. Zhang, and D. Wang, "Leveraging energy storage to optimize data center electricity cost in emerging power markets," *Proceedings of the 7th International Conference on Future Energy Systems, e-Energy 2016*, Jun. 2016, 10.1145/2934328.2934346, /doi/pdf/10.1145/2934328.2934346?download=true.

[48] F. Núñez, D. Canca, and A. Arcos-Vargas, "An assessment of European electricity arbitrage using storage systems," 2022, 10.1016/j.energy.2021.122916, https://doi.org/10.1016/j.energy.2021.122916.

[49] Á. Arcos-Vargas, D. Canca, and F. Núñez, "Impact of battery technological progress on electricity arbitrage: An application to the Iberian market," *Applied Energy*, vol. 260, Feb. 2020, ISSN: 03062619, 10.1016/j.apenergy.2019.114273, https://ideas.repec.org/a/eee/appene/v260y2020ics0306261919319609.html.

[50] L. Becchi et al., "A Computationally Efficient Rule-Based Scheduling Algorithm for Battery Energy Storage Systems," *Sustainability 2024, Vol. 16, Page 10313*, vol. 16, no. 23, p. 10 313, Nov. 2024, ISSN: 2071-1050, 10.3390/SU162310313, https://www.mdpi.com/2071-1050/16/23/10313/htm%20https://www.mdpi.com/2071-1050/16/23/10313.

[51] Modo Energy, *Cycling your battery: what's the value of a cycle?* 2023, https://modoenergy.com/research/en/battery-energy-storage-cycling-rates-value-wholesale-frequency-response.

[52] M. Karasavvidis, D. Papadaskalopoulos, and G. Strbac, "Optimal Offering of Energy Storage in Electricity Markets With Loop Blocks," *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 299–315, Jan. 2024, ISSN: 15580679, 10.1109/TPWRS.2023.3240079, https://ieeexplore.ieee.org/document/10026336.

[53] I. Savelli, B. Cornélusse, A. Giannitrapani, S. Paoletti, and A. Vicino, "A new approach to electricity market clearing with uniform purchase price and curtailable block orders," *Applied Energy*, vol. 226, pp. 618–630, Sep. 2018, ISSN: 03062619, 10.1016/j.apenergy.2018.06.003, http://arxiv.org/abs/1711.07731%20http://dx.doi.org/10.1016/j.apenergy.2018.06.003.

A | Appendix

A.1 | Pipeline Architecture and Reproducibility

Interactive Dashboards

Every figure, table, and statistic discussed in this appendix and in Chapter 5 can be explored interactively, at the level of individual days and units, in one of three self-contained HTML dashboards built directly from the pipeline's outputs:

- **Market Overview & Backcast**
  unit-level day-ahead/intraday market reconstruction — Appendix A.1
- **Block-Order Submissions**
  full daily submission timeline and per-unit block-order log — Appendix A.2
- **BESS Backtest Results**
  full backtest sample across all three intraday sessions — Appendix A.4

The empirical and backtest results in Chapter 5 are produced by a sequence of independent Python scripts, each reading the previous stage's output and writing its own to disk, rather than a single monolithic program. This section documents that sequence at the level of what each stage consumes and produces, so that the process is auditable without requiring the underlying codebase. The pipeline separates into two halves: a *data extraction* stage that pulls and parses raw OMIE price files, and a *BESS modelling* stage that forecasts, dispatches, and backtests the battery against those prices.

Stage 1: Data extraction. Two download scripts retrieve historical OMIE price files directly from OMIE's public file-download endpoint, one per market:

- `1_download_prices.py` retrieves day-ahead `marginalpdbc` files, one per calendar day, for a specified date range, skipping any file already present locally.
- `download_prices_intraday.py` retrieves intraday `marginalpibc` files, one per session (IDA1–IDA3) per calendar day, over the same mechanism.

Both scripts write raw, unparsed OMIE files to local directories; no forecasting, optimisation, or aggregation happens at this stage. This is the only stage that depends on network access; every later stage operates entirely on the locally cached files.

Stage 2: Price forecasting (LEAR-LASSO). `2_price_forecast.py` and `2_price_forecast_ida.py` implement the LEAR model described in Section 4.5 and summarised in Appendix A.3, for the day-ahead and intraday markets respectively. Each parses the raw price files from Stage 1, constructs the 17-feature matrix per quarter-hourly slot, and exposes the fitted models as a library (`load_all_prices`, `fit_models`, `predict_day`) consumed directly by the optimisation stage rather than writing forecasts to disk independently; the intraday variant fits one such model per IDA session, with Session 3 predicted only over its own P49–P96 window.

Stage 3: Order-type optimisation and backtest. `3_order_optimizer.py` (day-ahead) and `6_intraday_optimizer.py` (intraday) implement the battery LP dispatch models for Simple, Loop, and the perfect-foresight oracle (Section 4.6), consuming the Stage 2 forecaster as an imported module and solving each day's dispatch with `scipy`.

`optimize.linprog`. Run in backtest mode, `6_intraday_optimizer.py` writes `intraday_backtest.json`: one record per calendar day, containing each session's realised and forecast revenue for every order type, the LP dispatch schedules, the raw forecast and actual price series, and the grid-search Loop specification. A companion welfare-only mode writes `intraday_welfare.json`, the perfect-foresight oracle computed independently of any forecaster. Both scripts checkpoint progress after every simulated day, so an interrupted run resumes rather than restarting.

Stage 4: Rule-based Loop rescheduling. `patch_rule_based_loop.py` reads `intraday_backtest.json` and replaces only the Loop order's revenue and window specification with the rule-based heuristic result described in Section 5.3.2; every other field (Simple, oracle, forecast, actual prices) is copied through unchanged. It writes `intraday_backtest_ruleloop.json`, the file underlying every Loop-related figure and table in Chapter 5. This separation keeps the two Loop scheduling methods directly comparable against an identical Simple/oracle baseline, since neither the LEAR forecast nor the LP dispatch is re-run between them.

Stage 5: Figures and reporting. `generate_summary_figures.py` and `generate_dispatch_figures.py` read `intraday_backtest_ruleloop.json` and write the static PNG figures included throughout Chapter 5 (win-share, revenue distributions, gap histograms, and per-day dispatch charts). `5_build_backtest_report.py` and `7_build_intraday_report.py` independently assemble self-contained interactive HTML dashboards (Chart.js-based) from the same JSON outputs, used during development to inspect individual days rather than as thesis exhibits.

Reproducing a result. Any headline number in Chapter 5 can be traced back through this chain: a figure or table cites a session and date, which can be looked up directly in `intraday_backtest_ruleloop.json` (or `intraday_backtest.json` for the grid-search comparison), which was produced deterministically from `marginalpibc` files retrievable from Stage 1 given the same date range. No stage depends on manual intervention beyond invoking the scripts in sequence.

The resulting HTML file, `backcast_unified.html`, of the order of a few megabytes for a full day including intraday, exposes four tabs that together reconstruct the market outcome at unit level:

1. **Market Overview**: clearing-price curve, top-15 revenue decomposition into inframarginal rent and residual, and a searchable, sortable, paginated table of all matched units with a sparkline profile and a drill-down panel for each. For 1 December 2025 the dashboard reports 457 GWh of matched Spanish generation, €41.2 M of cleared revenue, a generation-weighted price of 90.25 €/MWh, and a day spread of 96.36 €/MWh between the cheapest and the most expensive period.
2. **Intraday Market**: a per-session view with price comparison against day-ahead, incremental revenue KPIs computed against the previous session's baseline, top-15 revenue chart, and the unit-level table restricted to the session's re-auctioned periods and units. On the reference day, the three sessions move 169.8, 29.1 and 10.6 GWh of incremental volume and 14.2, 2.4 and 0.8 M€ of incremental revenue respectively, with the sharp decay illustrating the well-documented falling liquidity of late intraday sessions.

3. **Data Integrity Check**: the pumped-hydro revenue cross-check described above, with a selectable plant and its pump/turbine/price profile overlay.
4. **Complex Order Compliance**: per-unit SCO cards showing income condition and MAV compliance, together with a summary table classifying every complex-offer unit as activated, withdrawn, or anomalous.

A.2 | Data Provenance and Sample Construction

Raw OMIE files. Two OMIE file families underpin the intraday backtest and the block-order study. The findings in this section are drawn from and can be explored interactively in `multimonth_blocks.html`, which exposes the full daily submission timeline, technology and session breakdowns, clearing-rate distributions, and a per-unit block-order log for the complete sample.

- `marginalpibc`: the IDA auction marginal (system) clearing price series, one file per session per day, used as both the LEAR forecasting target and the realised-price series against which Simple, Loop, and the oracle LP are settled.
- `curva_pibc_uof`: unit-level offer curves for the IDA auctions, used exclusively for the block-order usage study (Section 5.2): every submitted simple and block offer, by unit, technology, direction, price, and declared period span.

The equivalent day-ahead file, `marginalpdbc`, is used only for the day-ahead cross-check (Section 5.3.13) and is not part of the intraday backtest proper.

Sample windows. The block-order study covers 1 February 2025 – 31 January 2026 (365 calendar days). The intraday BESS backtest covers a separate, later-closing window, 5 February 2025 – 11 May 2026 (461 calendar days), of which 433 have a usable price file for every session; the remaining 28 days are excluded outright for missing or malformed source files, prior to any burn-in filtering.

Resolution transition. OMIE moved from 24-period hourly to 96-period quarter-hourly intraday scheduling in October 2025. Pre-transition files are upsampled $\times 4$ (each hourly price repeated across its four constituent quarter-hours) so that every session in the sample is represented on a consistent 96-slot grid, both for LEAR's per-period estimators and for the LP dispatch models.

Burn-in exclusion. The LEAR model requires an expanding training window before its first forecast is well-specified (Appendix A.3): 21 days of minimum training data plus a 14-day longest autoregressive lag, giving a 35-day burn-in per session. Of the 433 days with usable price files, this leaves 398 post-burn-in days for Session 1, 395 for Session 2, and 393 for Session 3 (Session 3's slightly different count reflects its own independent LEAR model and burn-in calculation over the P49–P96 window). All headline figures, tables, and per-session statistics in Chapter 5 are computed on these post-burn-in samples; days within the burn-in window are excluded rather than scored, since the model has no forecast to act on and would otherwise contribute a structural, uninformative zero to every metric.

The "combined" 387-day sample. Where Sessions 1, 2, and 3 are compared or summed on the same calendar day (e.g. Table 5.8), the sample is further restricted to the 387 days on which *all three* sessions have simultaneously completed their own burn-in. This is a strict subset of each session's own 398/395/393-day sample and is not their arithmetic sum; it is the intersection of three independently-determined post-burn-in day sets.

A.3 | LEAR-LASSO Specification Summary

Section 4.5 describes the LEAR forecasting model in prose; Table A.1 collects the same specification in a single reference table for reproducibility.

Table A.1: LEAR-LASSO model specification.

Parameter	**Value / description**
Model family	LASSO-regularised linear regression (Lasso Estimated AutoRegressive, LEAR)
Estimators per session	96 independent models, one per quarter-hourly period; each trained only on that period's own historical prices
Regularisation	ℓ_1 (LASSO); penalty strength α selected by cross-validation (LassoCV) on the training window at each retraining step
Feature count	17
Same-slot lags (6)	$p_{d-1,t}$, $p_{d-2,t}$, $p_{d-3,t}$, $p_{d-7,t}$, $p_{d-14,t}$, and the 7-day same-slot rolling mean
Daily-level (5)	$\bar{p}_{d-1}$, $\bar{p}_{d-2}$, $\bar{p}_{d-7}$ (daily means), plus yesterday's within-day min and max
Calendar (6)	sine/cosine pairs for day-of-week, month-of-year, and intra-day period
Training scheme	Expanding walk-forward (rolling-origin); no historical observations discarded
Burn-in / first forecast	Day 35 of the sample (21-day minimum training buffer + 14-day longest lag)
Retraining cadence	Daily; a fresh set of 96 models is re-estimated for every forecast day, using only data available before that day's session closes
Per-session independence	A separate LEAR model is fit for each of Sessions 1, 2, and 3, targeting that session's own realised `marginalpibc` series; no cross-session features are used
Pre-Oct-2025 resolution	24 hourly prices upsampled ×4 to populate the 96-slot quarter-hourly grid
Evaluation metric	MAE (mean absolute error, €/MWh), averaged across all periods and evaluation days

BESS backtest dashboard. The companion battery-results file, `intraday_report_ruleloop.html`, built by `7_build_intraday_report.py` from `intraday_backtest_ruleloop.json`, exposes tabs covering the full backtest sample rather than a single reference day:

1. **Session 1, Session 2, and Session 3 results**: one tab per session, each with an order-type selection summary (share of days won by Simple, Loop, and NoBid), a daily revenue chart comparing Simple, Loop, and the oracle, a monthly average revenue chart, a histogram of the gap to oracle, a Simple-versus-Loop revenue scatter, and a full clickable day-by-day table that expands into a per-day price chart on selection.
2. **Cross-session comparison**: the same summary views aggregated across all three sessions, giving the combined-day figures reported throughout Chapter 5 (e.g. the 387-day combined sample, the 81.4 % oracle capture) in an interactive form that can be re-filtered by date range or order type.

Both dashboards read directly from the JSON outputs already documented in Stages 3–4 above, so regenerating either one after a pipeline re-run requires no additional data preparation beyond re-invoking the corresponding build script.

A.4 | Loop Clearing: Worked Numerical Examples

Section 5.3.8 discusses the C88 clearing condition and several representative days at a summary level. This appendix works through the underlying leg-level arithmetic for two additional days, showing exactly how the declared windows, realised prices, and the net-margin condition $\lambda_{h_2} - \lambda_{h_1} \geq c_{\text{deg}}/\eta$ combine to produce the reported revenue.

Example A: 30 October 2025, Session 1 (two-cycle, one leg rejected). The rule-based search proposes two cycles. Loop1 charges 8.75 MWh over 13:30–15:15 at an average realised price of 19.2 €/MWh (total cost 168.0 €) and discharges 6.25 MWh over 18:45–20:00 at 112.9 €/MWh (total revenue 705.5 €); after 62.5 € of degradation (6.25 MWh × 10 €/MWh), this leg nets 475.0 €. Loop2 charges 3.75 MWh over 04:15–05:00 at 89.9 €/MWh (337.1 €) and discharges 2.50 MWh over 09:00–09:30 at 103.9 €/MWh (259.7 €); after 25.0 € of degradation, this leg's realised net margin is $259.7 - 337.1 - 25.0 = -102.4$ €, which fails the clearing condition and is recorded as 0 € (Section 5.3.11). The session total is therefore $475.0 + 0.0 = 475.0$ € on a forecast basis; the corrected realised figure of 372.6 € reported in Table 5.12 instead credits Loop2's true realised loss of -102.4 € directly, rather than flooring it to zero. Simple's own realised revenue this session is 283.5 €, against an oracle of 799.2 €.

Example B: 21 October 2025, Session 1 (a well-priced single leg against a strongly fragmented Simple schedule). The search proposes one cycle: charge 8.75 MWh over 14:45–16:30 at an average realised price of -3.0 €/MWh (total cost -26.6 €, i.e. the battery is paid to charge into this window) and discharge 6.25 MWh over 19:30–20:45 at 91.3 €/MWh (total revenue 570.9 €); after 62.5 € of degradation, the leg clears at $570.9 - (-26.6) - 62.5 = 535.0$ €, matching the realised total of 534.96 € reported in Section 5.3.8 up to rounding. Simple's LP, over the same session, dispatches ten separate charge and discharge blocks spanning an overnight cycle (02:30–05:15 charge, 07:45–09:30 discharge) and a second afternoon/evening cycle overlapping Loop's declared window, reaching 1,233.7 € against an oracle of 1,251.9 €. This pair illustrates directly why Loop's single best-scoring leg, even when correctly identified and profitably cleared, cannot match a schedule that legitimately requires two separate cycles across the session.

A.4.1 | Manual to read the BESS model graphs

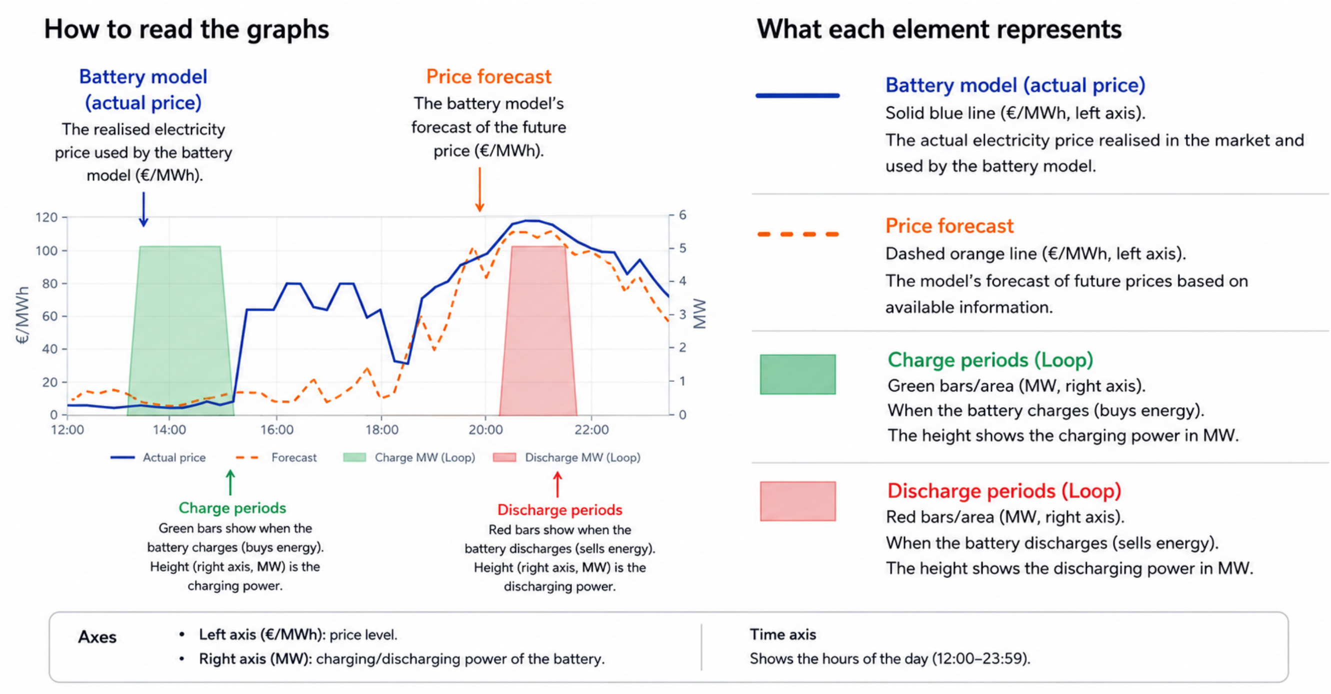

Figure A.1: Manual to read the graphs (from the model and the html