EuroSun2026

Cascading Machine Learning for Sub-Hourly BESS Dispatch: Paving the Way for Hybrid PV-Storage Integration

Noam Benazouz, Farzad Ghasemi, Thomas Carrière, Sébastien Pitaval, Christophe Vernay, Nicolas Thévenin

CalibSun, Solaïs, Sophia Antipolis, France

noam.benazouz@calibsun.com

Summary

This study evaluates the impact of production forecasts on the revenues of a hybrid power plant. The system combines a 5 MWp solar plant and a 2.27 MW / 4.55 MWh battery. Both share a grid connection in northern Spain. The plant is simulated over 243 days on the Iberian 15-minute electricity market. It submits bids in four daily auctions and dispatches the battery every 15 minutes. The simulation is executed four times to compare different production forecasts. All runs use identical price models, bidding strategies, control algorithms, and battery degradation costs. The evaluated inputs are a site-calibrated commercial forecast, the GFS weather model, a smart persistence baseline, and the actual measured irradiance. This measured irradiance represents a perfect forecast and defines the maximum theoretical revenue. The plant earns 152.5 kEUR with smart persistence and 178.1 kEUR with perfect production. This creates a theoretical gap of 25.5 kEUR. The commercial forecast generates 161.9 kEUR, which captures 36.5 % of this gap. The GFS model earns 158.2 kEUR, capturing 22.2 %. The commercial forecast also yields a mean absolute error 10 % lower than GFS. Battery usage remains stable across the three realistic cases, totaling about 346 equivalent full cycles. However, better forecasting reduces the energy deviation from the trading plan. This deviation falls from 1.47 GWh with smart persistence to 1.22 GWh with the commercial forecast. Furthermore, a model predictive controller (MPC) is compared to a simple tracker that strictly follows the bidding plan. The MPC increases revenues by 7.1 % when production is known in advance. When production is uncertain, it increases revenues by 18 to 20 %.

PV production forecasting, forecast value, battery storage, hybrid plant, 15-minute market, imbalance settlement, mixed-integer optimisation, LightGBM

Field of application for the findings: operation and revenue assessment of utility-scale hybrid PV–storage plants; valuation of commercial PV forecasting products.

Addressed audience: plant operators and asset managers, forecast providers, market-design analysts.

1. Introduction

High photovoltaic (PV) penetration has changed the price dynamics for solar power plants. On the Iberian market, midday hours are now typically the cheapest and evening prices rise significantly as solar generation decreases. Furthermore, the market time resolution has shifted from hourly to quarter-hourly intervals. This transition occurred on 19 March 2025 for the imbalance settlement period, and on 1 October 2025 for the day-ahead market. A hybrid plant combining PV and battery storage under a single grid connection must make two distinct decisions: bidding energy volumes on the market in advance, and physically dispatching power in real time. Any deviation between the cleared bids and the actual energy delivery is settled at the imbalance prices presented in Tab. 1.

The literature on hybrid PV–storage dispatch is large. Studies model bidding across day-ahead and intraday sessions, and they add imbalance settlement rules to penalise deviations. Solving the problem requires three components. The first is a price forecast. Electricity price forecasting is a mature field, and its methods and benchmarks are written for the day-ahead auction (Weron, 2014; Lago et al., 2021). The second component is a battery degradation model. The third is a control strategy, usually model predictive control or stochastic programming. The imbalance signal is treated apart from these three. Published work forecasts it at horizons of minutes rather than of a day (Bottieau et al., 2020).

Most of these studies focus on the control strategy. They treat the production forecast as fixed data, or they assume its error follows a known law. Carrière (2020) couples probabilistic PV forecasts to a market policy and analyses the profitability of a PV–battery system. The forecast product is not itself the variable there. Studies that compare production forecasts against each other are rare.

This paper measures the value of a production forecast, on one plant and one market. It makes three contributions.

- A method to assign a financial value to a production forecast. A complete simulation chain models the process from the weather data to the final financial settlement. The simulation is non-anticipative, meaning it does not use future information. By changing only the production forecast between simulation runs, any difference in total revenue is strictly caused by the forecast quality.
- A measured hierarchy of four production forecasts, reported as revenue of the whole plant over the whole window, with and without the permission to charge the battery from the grid.
- A design rule for real-time plant controllers. When operating under imbalance pricing, the controller must apply a penalty for deviating from the initial market bid. This internal penalty must be calibrated to account for the large estimation errors associated with imbalance prices.

2. Market and asset

2.1 The gate sequence

On the Iberian market, energy is traded through a series of auctions. Since June 2024, the Single Intraday Coupling (SIDC) has replaced the national intraday sessions with three pan-European intraday auctions. There are therefore four trading windows: one day-ahead and three intraday auctions. Every session clears at a quarter-hourly resolution over the whole study window. The day-ahead closes at 12:00 on day D-1 (Madrid time). The three intraday auctions close at 15:00 and 22:00 on D-1, and 10:00 on D. The first three cover all 96 slots of day D. The third auction only covers 12:00 to 24:00. Once a window closes, the volume bid is final and binding.

2.2 Dual imbalance pricing

In the Iberian market, energy delivered differs from energy nominated, and this mismatch is settled at two asymmetric prices: one for undersupply and one for oversupply. Every plant is represented on the market by a balance-responsible party (BRP), which is the entity held accountable for the gap between what it nominated and what it delivered. The two prices differ because the system operator pays for the balancing energy it activates. A BRP that deviates in the direction the system needs is paid less than the market price; one that deviates against it pays more.

A BRP that delivers less than its promised volume is "short" and pays the deficit price π^- on the difference. One that delivers more is "long" and is paid the surplus price π^+. The two terms are the trading vocabulary rather than settled scientific usage, and they are used here for brevity.

Take one quarter-hour. Write R for the cash it settles, in euros: the money received for the energy sold, plus or minus what the deviation is worth. Write n for the volume sold and q for the energy delivered, both in MWh. Write π for the price of the session the volume was sold in, day-ahead or intraday. Write $\pi^\pm$ for the price that settles the deviation, that is π^- when the plant is short and π^+ when it is long. The cash can be written on the volume sold,

$$R = \pi\, n + \pi^\pm (q - n), \tag{1}$$

or on the energy delivered,

$$R = \pi\, q + \left(\pi - \pi^\pm\right)(n - q). \tag{2}$$

The two forms are equal. The second one prices every MWh of deviation at the coefficient $\pi - \pi^{\pm}$, and that coefficient is the deviation margin. Taking the two cases apart gives

$$c^{-} = \pi^{-} - \pi, \qquad c^{+} = \pi - \pi^{+}, \tag{3}$$

the cost of being one MWh short and one MWh long. Eq. 2 then reads $R = \pi\, q - c^{-} s - c^{+} \ell$, with $s = (n-q)^{+}$ the short leg and $\ell = (q-n)^{+}$ the long leg. A superscript on π always marks an imbalance price; π on its own is the price of a market session. All these quantities are defined per quarter-hour; the slot index is dropped here and restored in Section 3.3.

The deficit and surplus prices (π^{-} and π^{+}) are set by the transmission system operator (TSO) from the balancing energy it actually activated. The two mechanisms are therefore independent. Over the window π^{-} or π^{+} equals the day-ahead price on 0.16 % of slots, and nothing forces either of them above or below the price of a session. The ordering $\pi^{-} \geq \pi^{+}$ holds on 93.6 % of the 23 328 slots where all three prices are published, and it fails on the remaining 6.4 %. A deviation margin can therefore be negative, which means that deviating can pay. Tab. 1 gives the statistics.

Tab. 1: Imbalance pricing statistics over the study window.

Metric	Value
Short leg margin negative ($c^{-} < 0$)	32.9% of slots
Long leg margin negative ($c^{+} < 0$)	22.7% of slots
Deficit price premium over day-ahead	13.4 EUR/MWh
Surplus price discount to day-ahead	17.2 EUR/MWh
Peak hours (18:00–21:00), short pays	42.0% of slots
Off-peak hours, short pays	31.6% of slots

2.3 The plant

The test site is a reference radiometric station in Navarre, northern Spain, operated by the Centro Nacional de Energias Renovables (CENER). It measures global horizontal irradiance and ambient temperature, and a commercial provider issues operational irradiance forecasts for it (Section 3.1).

The plant is not a metered asset. It is the documented configuration of a 5 MWp single-axis-tracking plant, paired here with a 2.27 MW / 4.55 MWh battery (Tab. 2). The plant is modelled and not measured. In Spain, the solar plant and the battery share one grid connection and one export limit. The export limit is 4.55 MW for 5.0 MWp of panels, so the plant clips when irradiance is high. The energy above 4.55 MW can be stored in the battery instead of being lost. The battery can also buy power from the grid.

Production is calculated from measured irradiance, not from meters. The same solar model is used for the actual irradiance and for all four forecasts. No forecast can therefore differ from the truth because of the model.

Over 243 days, the plant produced 790.6 kWh per kWp. In perfect clear sky, it would produce 1 327.1 kWh per kWp. The average clear-sky index over the window is therefore 0.596, and it drops to 0.42 in January. This is a cloudy site compared to the Mediterranean coast, and prediction is harder here. The grid export limit is reached during only 0.4 % of the daylight intervals. Therefore, storing clipped energy provides a marginal benefit.

3. Method

The method has four parts and they run in the following order. First, four production forecasting models produce an irradiance forecast each (Section 3.1). Second, a cascading machine-learning model forecasts the price of every session (Section 3.2). Third, a model estimates the cost of deviating from the nomination (Section 3.3). Fourth, two decision layers use those three inputs: a mixed-integer program bids at each gate (Section 3.4) and an MPC controller decides the physical dispatch every quarter hour (Section 3.5). Only the first part changes between the four runs.

3.1 Four production forecasting models

Four models are compared. They share the same asset, time window, price forecasts, two decision layers, and transposition model. The irradiance forecast is the only difference between them.

Tab. 2: Plant, battery and window. [a] POI, point of connection: the single grid access point the PV and the battery share. [b] EFC, equivalent full cycle: one full charge and discharge of the rated energy. [c] LFP, lithium iron phosphate. The two grid-charging flags are coupled: with imports forbidden, the non-negativity of the injection already forces the battery to absorb no more than the instantaneous PV.

Quantity	Value	Quantity	Value
PV peak power (DC)	5.0 MWp	Battery power	2.27 MW
Export limit at the POI[a]	4.55 MW	Battery energy	4.55 MWh
DC/AC ratio	1.10	Duration	2.0 h
Tracking	single-axis	Chemistry	LFP[c]
Backtracking	yes	One-way efficiency	0.955
Row pitch / module length	4.0 / 2.0 m	Round-trip efficiency	0.912
Maximum tracker angle	55°	Marginal degradation cost	3 EUR/MWh
Performance ratio	0.85	State-of-charge window	5–95 %
Temperature coefficient	$0.004\ \mathrm{K}^{-1}$	Warranty cycle cap	1.5 EFC/day[b]
Latitude / longitude	42.816 / −1.601	Trading fee	0.25 EUR/MWh
Elevation	471 m	Settlement slot length	15 min
Study window	2025-10-01 – 2026-05-31	Delivery days	243

NEXT is a commercial irradiance forecast calibrated on the site's own measurements and published as quantiles. Calibrating on the site could let the model see its own future, so the calibration is walk-forward. Each run comes from a rolling backtest that refits only on measurements published strictly before its issue time. No model is fitted on a day it then predicts.

GFS is the Global Forecast System of the National Oceanic and Atmospheric Administration, a global numerical weather prediction used here as a reference.

Smart persistence assumes the atmosphere does not change over the horizon. It computes the clear-sky index $k^* = G/G_{\mathrm{cls}}$, where G is the measured global horizontal irradiance and G_{cls} its geometric clear-sky value. The index is capped at 1.1, held constant over the horizon and converted back to irradiance.

Perfect irradiance is the realised measurement. It is used as an upper bound and not as a forecast. What a model loses relative to it comes from forecast error.

Each of the three, and the realised irradiance itself, is then converted to power by the single transposition of Section 2.3. The PV model is therefore a common factor and not a per-model choice. Its own error does not enter the comparison, because the same model produces the truth. The perfect run is the measured irradiance passed through that same transposition, so its production is exact. The study measures the value of an irradiance forecast on this plant. It does not measure the value of a PV model.

3.2 Cascading quantile forecasts of the session prices

The gate program must bid against prices it does not know. Prior to the day-ahead closure at 12:00 D−1, the model generates absolute price forecasts to establish a baseline schedule. Forecasting the intraday sessions as absolute prices as well would propagate that error into every later gate, whereas in operations the day-ahead prices clear before IDA1 closes and can serve as a deterministic baseline. A hierarchical cascading architecture is therefore used (Fig. 1): rather than forecasting absolute intraday prices, each subsequent model predicts only the price differential, or spread, between the known anchor and the upcoming session.

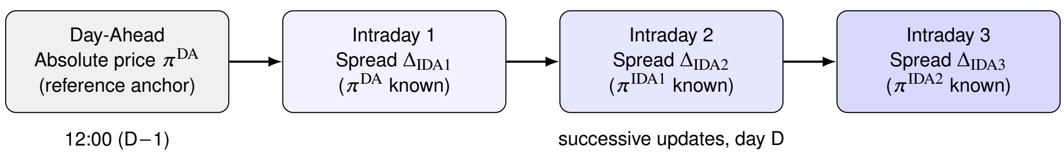

Fig. 1: Cascading forecasting timeline: from an absolute day-ahead baseline to iterative spread corrections. Each session's spread is reconstructed on the realised price of its anchor, so no gate uses information published after it closes.

Among the methods tested, quantile regression on LightGBM (Ke et al., 2017) is selected. Two properties motivate the choice. Gradient boosting fits shallow trees in sequence, each one learning from the errors of the previous one. It also handles heterogeneous tabular inputs without scaling them. The median forecast ($\tau = 0.5$) is used, because it is less affected by the extreme prices this market produces.

One model is fitted per target, so four in total: one for the day-ahead price and one for each intraday spread.

All four share the same feature set: system load forecast, wind and solar forecasts, local hour, day of week, and past prices. The feature set excludes price and spread columns, so the anchor enters through the reconstruction only.

The forecast never looks ahead in time. The day-ahead auction closes before the first intraday window opens. Each window only uses information available when it closes.

3.3 Deviation margins

The bidding and dispatch steps never use raw imbalance prices. Instead, both rely on the two imbalance margins defined in Eq. 3. These margins are the actual quantities being forecasted. They are forecasted directly without truncating negative values. A negative margin indicates that the market settlement rewards deviations, and both decision steps are designed to process this financial opportunity.

The bidding and dispatch steps require different estimators due to distinct forecasting horizons. At the day-ahead bidding gates, decisions are made 12 to 36 hours in advance. At this stage, recent imbalance prices are unavailable because they are published only after a time slot concludes. The time of day is the sole available predictor. However, this variable explains only 3.3 % of the variance for the short margin and 3.8 % for the long margin. Over 96 % of the variance remains unexplained. Consequently, the estimator for the bidding gate relies on a simple 30-day rolling mean for each specific time slot.

Using a rolling mean successfully captures the structural asymmetry between the two imbalance prices, as shown in Tab. 1. This asymmetry allows power plants to optimize market bids without requiring precise short-term price forecasts. Pinson et al. (2007) show this for wind generation: the revenue-maximising bid is not the expected production but a quantile of its predictive distribution, and the level of that quantile follows from the ratio of the two regulation costs. The present hybrid model applies this principle by integrating the historical price asymmetry directly into the optimization objective through the variables $\hat{c}^-$ and $\hat{c}^+$. Consequently, the algorithm simply trades off these asymmetric market penalties against the operational cost of cycling the battery.

At dispatch the estimate becomes usable. The MPC solves a new problem every quarter hour, and the margin of the previous slot is known by then. The imbalance price is derived from the balancing energy the TSO has already activated, so it is published about fifteen minutes after the slot it settles rather than a day later. An autoregressive term is therefore available. The estimator blends the last known margin with the average margin for the same time of day,

$$\hat{c}_{t+h} = \alpha_h \, c_{t-1} + (1 - \alpha_h) \, \mu_{k(t+h)}, \tag{4}$$

where h is the lead in slots, $k(\cdot)$ maps a slot to its time of day, and μ_k is the average realised margin over the 30 preceding days at that time of day. One forecast therefore covers the whole horizon of a single optimisation, with a weight that decays from the near slots to the far ones.

The weight α_h is empirically fitted rather than derived from theoretical autocorrelation. For each forecast horizon and each imbalance margin, α_h is tested on a grid from 0 to 1 in steps of 0.05. The algorithm selects the value that minimizes the mean absolute error over the preceding 30 days. For a one-slot horizon, the optimal weight is 0.80. This weight decreases as the forecast horizon extends, dropping to 0.50 at half an hour, 0.25 at two hours, and 0.10 beyond. The estimator has two inputs and neither of them wins everywhere. The last known margin is the better of the two close to delivery. Past roughly one hour it becomes the worse of the two, and the time-of-day average takes over. The blend is more accurate than both inputs at every horizon tested, which is why the two are weighted against each other rather than swapped at a threshold. Section 4.1 reports the errors.

3.4 The gate program

At each gate, an optimization problem is solved. It chooses how much to bid over the slots still open and when to charge or discharge the battery, and it takes the volumes bid at earlier gates as sunk.

The program uses the second form of the settlement identity, Eq. 2, and not the first. The two forms are algebraically equal, but they do not behave the same inside an optimisation. Written on the volume sold, the nomination n is free to run away from the injection whenever π^- sits below the session price, and the program is unbounded. Written on the energy delivered, the deviation of a slot is bounded by what the plant can physically inject, which is what makes the program solvable.

The objective maximised over the open slots is

$$\sum_t \Delta t \Big[(q_t - n_t^{g-1})\, \pi_t^g - c_{\text{tr}} v_t - \kappa_t^- s_t - \kappa_t^+ \ell_t - c_{\text{deg}} (c_t + d_t) \Big] + \lambda_T e_T, \tag{5}$$

with $v_t = |n_t - n_t^{g-1}|$ the volume this gate trades, which is what the fee applies to, c_t and d_t the battery charge and discharge, e_T the state of charge at the horizon, and deviation coefficients

$$\kappa_t^- = \hat{c}_t^- + \delta + \varepsilon, \qquad \kappa_t^+ = \hat{c}_t^+ + \delta + \varepsilon. \tag{6}$$

The term δ raises the price of deviating above the estimated margin. It is needed because $\hat{c}$ is an estimate and not a measurement. Section 3.5 explains what it protects against and how its value is set. The constraints are, for every open slot t,

$$\begin{aligned}
&q_t = p_t - w_t + d_t - c_t, && 0 \le w_t \le p_t, && \underline{Q} \le q_t \le \bar{Q}, \\
&e_t = e_{t-1} + \Delta t\, (\eta_c c_t - d_t / \eta_d)\,, && \underline{E} \le e_t \le \bar{E}, && c_t \le \bar{P} u_t, \\
&s_t \ge n_t - q_t, && \ell_t \ge q_t - n_t, && d_t \le \bar{P}(1 - u_t), \quad u_t \in \{0, 1\}.
\end{aligned} \tag{7}$$

together with the daily cycle cap $\Delta t \sum_t (c_t + d_t) \le 2\bar{E}N$, decremented by the cycling already spent that day. Here p_t is the forecast PV power and w_t the part of it that is clipped instead of injected. The balance needs w_t because PV above the export limit that the battery cannot absorb has nowhere to go, as in Section 2.3. $\bar{Q}$ is the export limit and $\underline{Q}$ its import counterpart, zero when grid charging is forbidden and $-\bar{Q}$ when it is allowed. η_c and η_d are the one-way efficiencies, $\bar{P}$ the battery power, $[\underline{E}, \bar{E}]$ the state-of-charge window and N the warranty cycle cap of Tab. 2. Degradation enters as a linear charge on throughput in the objective, on top of that cap.

Tab. 3: Notation of the gate program. Powers are in MW, energies in MWh, prices and coefficients in EUR/MWh, and Δt is the slot length in hours. The dispatch program of Section 3.5 reuses the same symbols with the nomination as data rather than as a variable. The term $\lambda_T e_T$ is required because the horizon ends at midnight: without a value on the energy left in the battery, the program would empty it at the end of every day whatever the next day is worth.

Symbol	Meaning	Symbol	Meaning
n_t	cumulative nomination	π_t^g	price of the session bid into
n_t^{g-1}	nomination sunk at earlier gates	$\hat{c}_t^-, \hat{c}_t^+$	estimated deviation margins
q_t	physical injection at the POI	κ_t^-, κ_t^+	deviation coefficients
$s_t = (n_t - q_t)^+$	short leg	δ	aversion to deviating
$\ell_t = (q_t - n_t)^+$	long leg	ε	tie-break, numerical
v_t	volume traded at this gate	c_{tr}	trading fee
c_t, d_t	battery charge, discharge	c_{deg}	marginal degradation cost
e_T	state of charge at the horizon	λ_T	terminal water value

Three points on Eqs. 5–6. First, the estimated margins $\hat{c}$ are not clipped at zero. Where the settlement rewards deviating, the coefficient goes negative and the solver is allowed to see the gain. That freedom has a cost in the formulation. The short leg is declared by two inequalities, $s_t \ge n_t - q_t$ and $s_t \ge 0$, which put a floor under s_t but no ceiling above it. No ceiling is normally needed: the objective charges κ_t^- for every MWh of s_t, so it pushes the leg down onto its floor and the leg lands on $\max(n_t - q_t, 0)$ by itself. A negative κ_t^- reverses that. The objective now pays for s_t, nothing above holds it back, and the program reports an arbitrarily large gain on a deviation the plant never makes. Each leg whose coefficient can turn negative is therefore forced to equal $\max(n_t - q_t, 0)$ exactly, by a binary variable and the physical bound on the deviation. Only the slots that need one receive a binary, and they are known before the solve: the coefficient is $\hat{c} + \delta + \varepsilon$, all three terms are data, so its sign is read off the margin estimate.

Second, $\varepsilon = 0.5$ EUR/MWh is numerical hygiene and not economics, which is why it is kept apart from δ. Where a margin lands near zero the objective goes flat in the nomination and the trading fee would drag it to zero, so the plant would appear to stop bidding for a 0.25 EUR/MWh reason. Unlike δ, ε is fixed once and never calibrated, and it is reported as such.

Third, the binary variable u_t in Eq. 7 prevents the battery from charging and discharging simultaneously. This ensures the market schedule cannot buy and sell energy during the exact same quarter-hour. However, the dispatch program described in Section 3.5 omits this binary constraint. This mathematical relaxation is safe because simultaneous charging and discharging is strictly suboptimal in this model, executing both actions at the same time does not alter the net physical

3.5 Real-time dispatch, and the aversion to deviating

A slot is frozen as soon as the last gate that covers it has closed. For the slots of 00:00–12:00 that happens at 22:00 the day before, when IDA2 closes; for the slots of 12:00–24:00 it happens at 10:00 on the delivery day, when IDA3 closes. The dispatch layer therefore runs on the frozen slots from midnight onward, while IDA3 is still able to revise the afternoon. The two layers are interleaved in clock order for exactly this reason.

Once a slot is frozen, the remaining decision is physical. The plant can move the battery to match the bid, or pay the imbalance price. Two control strategies are tested. The tracker tries to match the bid exactly. The model predictive controller (MPC) solves a new problem every 15 minutes over the slots that remain, applies its first action, and rolls forward. It decides at each step whether moving the battery is worth more than paying the imbalance.

The margin forecast has an error of 20 to 26 EUR/MWh. The margin itself averages a few EUR/MWh. The MPC would therefore act on estimation noise if it were given no penalty for deviating. A parameter δ is added, which raises the cost of deviating by the same amount on both legs.

The value of δ is set by a sweep. The full 243-day window is run once per candidate value and the value that earns most is kept. Section 4.2 reports the sweep. It gives $\delta = 10$ EUR/MWh. The same value is applied at both layers, gate and dispatch, so the agent has the same risk attitude when it bids and when it dispatches. This is the configuration used for every result that follows.

3.6 Protocol

The four models are simulated day by day. The planning and dispatch phases are interleaved chronologically to replicate real control room operations. For instance, the IDA3 auction closes at 10:00 on the delivery day. Its bids are anchored to the actual measured state of charge of the battery, rather than the theoretical trajectory from the planning layer.

Financial settlement is always calculated using actual realized prices. Incremental volumes are settled at the clearing price of their specific auction session, while residual deviations are penalized using the true realized imbalance prices. Battery degradation is billed linearly based on the physical wear cost. These strict rules apply equally to all models. The simulation runs cover the same 243-day window and share identical price forecasts, margin estimators, and physical settlement rules. The production forecast remains the sole variable between the runs.

All reported results represent the total revenues for the entire hybrid plant over this period. Furthermore, training data and simulation data never overlap. The price models are retrained every 30 days using only published prices that strictly precede the predicted block, starting from 2021-01-01. Similarly, the margin estimators only process historical data available prior to the decision instant. Finally, the production forecasts are generated from rolling backtests, refitted strictly before their respective issue times. This non-anticipative design guarantees that no future information is ever used to train a model.

4. Results

4.1 Forecast quality

Forecast error is measured separately at each of the four gates, on the forecast issued at that gate's closure. The horizons follow from the gate calendar of Section 2.1. The day-ahead gate closes at 12:00 on D−1 and covers the whole of day D, so its leads run from 12 to 36 hours. IDA1 closes at 15:00 on D−1, from 9 to 33 hours. IDA2 closes at 22:00 on D−1, from 2 to 26 hours. IDA3 closes at 10:00 on D and covers 12:00 to 24:00 only, from 2 to 14 hours.

Write $\hat{y}_t$ for the forecast power of slot t, y_t for the measured power, $\bar{P}_{\mathrm{ac}}$ for the AC capacity of Tab. 2 and N for the number of scored slots. The two error scores are

$$\mathrm{nMAE} = \frac{1}{N\bar{P}_{\mathrm{ac}}} \sum_t \left|\hat{y}_t - y_t\right|, \qquad \mathrm{nRMSE} = \frac{1}{\bar{P}_{\mathrm{ac}}} \left[\frac{1}{N} \sum_t \left(\hat{y}_t - y_t\right)^2\right]^{1/2}, \tag{8}$$

both reported in per cent. The mean absolute error measures the typical size of the error and the root mean square error weights the large errors more heavily. Night slots are excluded from the sums: they are trivially perfect and would dilute every score. A slot is scored when the measured or the forecast power exceeds 1 % of $\bar{P}_{\mathrm{ac}}$, which leaves about 10 200 slots per gate. The bias is reported on a different denominator, $\sum_t (\hat{y}_t - y_t)/\sum_t y_t$, so it reads as a share of the energy actually produced.

At the day-ahead gate, where most of the position is set, NEXT has an nMAE of 10.6 %, GFS has 11.8 % and smart persistence has 18.1 %. The ordering holds at all four gates. NEXT stays between 10.2 and 10.6 % across them and GFS between 11.1 and 11.9 %, while smart persistence improves from 18.1 % at the day-ahead gate to 14.1 % at IDA3, where its last measurement is only two hours old. GFS forecasts systematically too low, with a bias of −5.6 % at the day-ahead gate, while NEXT is nearly unbiased at +1.6 %.

The price forecast has an MAE of 13 EUR/MWh on the day-ahead price, better than a naive “same day last week” baseline by 28.9 %. On the three intraday windows, the price forecast is about as accurate as just using the previous window’s price.

The margin estimator of Section 3.3 is much weaker. Over the 243 days its blend has an MAE of 20.9 EUR/MWh at a 15-minute lead on the short leg and 25.7 EUR/MWh at 24 hours. On the long leg the two figures are 21.8 and 24.7. The freshly published margin cuts the error by 19.5 % fifteen minutes ahead and by 1.0 % a day ahead. The two decision layers therefore use different estimators. The error is also larger than the quantity being estimated, which is why the MPC carries the penalty of Section 4.2.

4.2 *Calibrating the deviation aversion*

The whole window is run once per candidate value of δ, on the deployable configuration: real forecast, aversion at both layers, grid charging allowed. Revenue over the window rises from 157.6 kEUR at $\delta = 0$ to 161.9 at 10, then falls back to 161.3 at 15, 160.3 at 20, 158.6 at 30 and 157.5 at 40. The optimum is interior, and both sides of it cost money. Below it the MPC trades its own estimation error. Above it the MPC declines deviations that the settlement would have paid for. Battery throughput moves very little, from 345 to 347 equivalent full cycles, so the penalty on deviating is not a penalty on cycling.

4.3 *The forecast hierarchy*

Tab. 4: Revenue of the whole plant over the 243 days, in thousands of euros, under the two controllers and the two grid-charging configurations. The last two columns give the absolute deviation from the nomination and the battery cycling of the deployed configuration, that is the MPC with grid charging allowed. All four rows share the asset, the window, the price forecasts, the margin estimators, the degradation charge and δ = 10 EUR/MWh; only the production forecast differs.

	No grid charging		Grid charging			
Production forecast	Tracker kEUR	MPC kEUR	Tracker kEUR	MPC kEUR	Deviation GWh	Cycling EFC
Perfect production	152.3	159.9	166.2	178.1	0.60	343
NEXT	130.2	143.1	134.8	161.9	1.22	346
GFS	126.8	139.5	133.7	158.2	1.20	346
Smart persistence	117.4	129.9	128.0	152.5	1.47	347

Table 4 shows the main results. In the deployed configuration, the distance between perfect production and the naive baseline is 25.5 kEUR over 243 days. NEXT captures 36.5 % of that distance and GFS captures 22.2 %.

NEXT earns more than GFS on 61.7 % of the days. The same order holds in the four columns of the table, that is with and without grid charging and under both controllers.

The three real forecasts cycle the battery by the same amount, 346 equivalent full cycles over the window for NEXT and for GFS and 347 for smart persistence. The deviation differs. Smart persistence leaves 1.47 GWh of imbalance over the window and NEXT leaves 1.22 GWh, for 9.4 kEUR more revenue.

Grid charging reduces the distance between the forecasts. Without it the distance is 30.0 kEUR, and NEXT captures 44.1 % of it against 32.0 % for GFS. With it the distance falls to 25.5 kEUR. Grid charging adds revenue to all four rows. It adds 22.6 kEUR to smart persistence and 18.2 kEUR to perfect production, which is why the distance between them narrows.

The MPC earns more than the tracker in all cases. With perfect production it adds 11.8 kEUR. With the three real forecasts it adds 24 to 27 kEUR.

4.4 One example day

Figure 2 shows 6 April 2026. The day-ahead price runs from −6.0 to 54.4 EUR/MWh and is zero or negative on 36 quarter-hours, from 09:45 to 18:30. The MPC earns 140 EUR that day and the tracker loses 309.

The battery is moved four times. It charges 4.48 MWh before sunrise at 35.0 EUR/MWh and discharges 2.00 MWh into the morning. It then charges 4.28 MWh during the middle of the day at −5.9 EUR/MWh, which means the plant is paid to absorb its own production. It discharges 2.84 MWh into the evening peak at 46.9 EUR/MWh. It finally charges 2.26 MWh again after 22:00, at 36 EUR/MWh, for the following morning.

That last move is the one the tracker does not make. The MPC ends the day with 3.50 MWh stored against 1.31 MWh for the tracker, and the two trajectories differ by up to 3.48 MWh over a usable window of 4.09. Both are bound by the same cycle cap, which the MPC reaches exactly at 1.500 equivalent full cycles. Over the day the plant buys 8.77 MWh from the grid at an average 15.0 EUR/MWh, deviates by 2.42 MWh in total and pays 36 EUR for it.

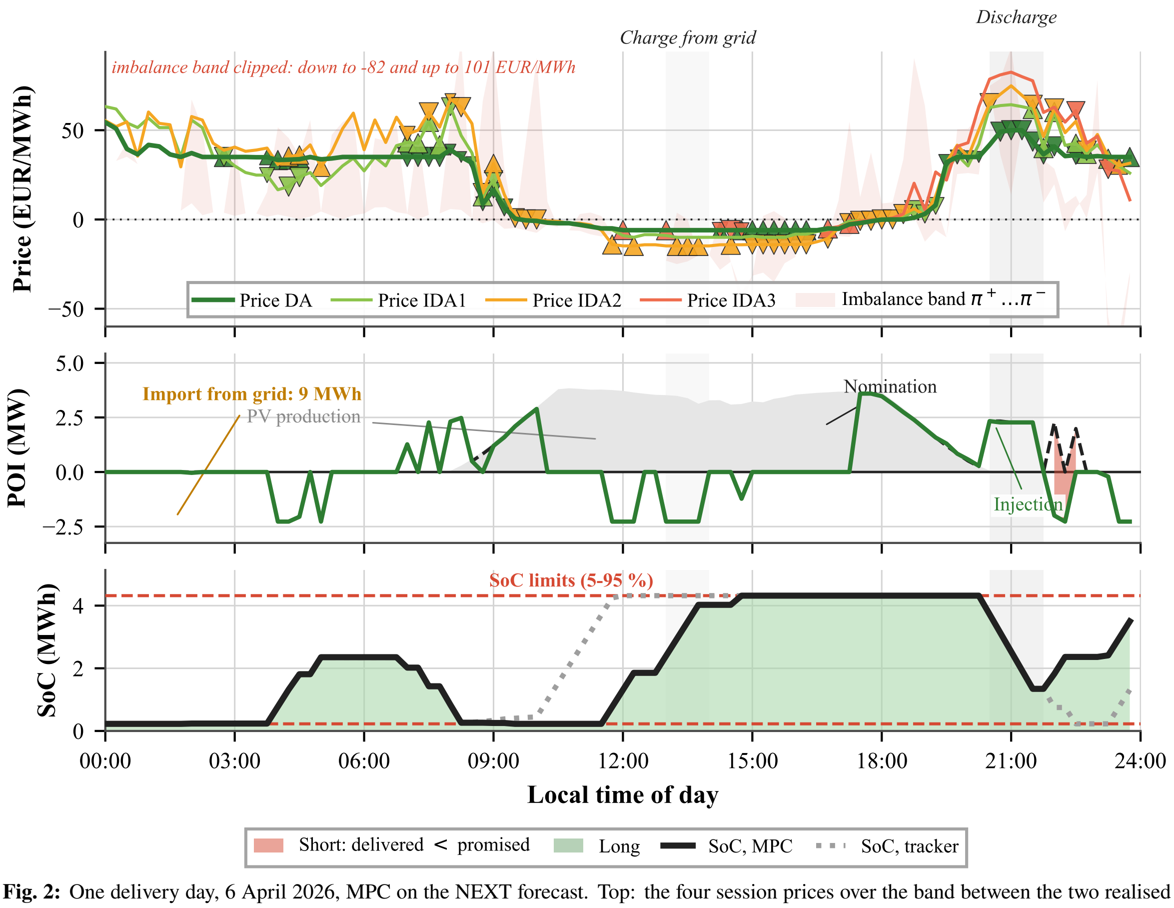

Fig. 2: One delivery day, 6 April 2026, MPC on the NEXT forecast. Top: the four session prices over the band between the two realised imbalance prices, each carrying the positions struck in that session, upward triangle to buy, downward to sell, marker area proportional to volume. Middle: promised versus delivered at the point of connection; the shaded gap is the imbalance, and negative injection is energy bought from the grid. Bottom: state of charge against its warranty window, the tracker dotted for contrast. Shaded bands label the operating regimes the dispatch produces.

5. Discussion

5.1 What bounds the value of a forecast and limitations.

There are two main factors constraining the financial value added by a production forecast. The first factor is the imbalance mechanism. Forecast errors induce financial penalties only when they translate into physical deviations. Providing the Model Predictive Control (MPC) with perfect margin forecasts increases revenues by 35 to 38 kEUR. This exceeds the 25.5 kEUR gap between the best and worst production forecasts. Therefore, the margin forecast accuracy limits the plant revenues more severely than the production forecast itself. The second factor is the market design. The final intraday auction, IDA3, closes at 10:00 for delivery between 12:00 and 24:00. Consequently, forecast improvements at shorter horizons cannot be monetized through market bids. They only inform the MPC, which is inherently constrained by the high unpredictability of real-time imbalance prices. Finally, the 243-day simulation window excludes the summer months, which typically feature the highest occurrence of negative prices. As a result, the reported revenues represent an upper bound of actual operational earnings.

6. Conclusion

Three production forecasts and a measured irradiance are tested on a hybrid PV–battery plant over 243 days on the Iberian 15-minute market. The same bidding and control system is used for all four, so only the forecast changes. The four forecasts rank in the same order in every configuration tested. The site-calibrated forecast captures 36.5 % of the distance between the naive baseline and perfect production, and the free weather model captures 22.2 %.

The MPC needs a penalty for deviating from its bid. The margin it trades is estimated with an error of 20 to 26 EUR/MWh, for a quantity that averages a few EUR/MWh. Without the penalty the MPC acts on that error. The penalty is calibrated on revenue and it is worth 4.3 kEUR over the window.

Two results concern operation. Grid charging adds 18 to 23 kEUR to every configuration, and it adds most to the weakest forecast. Real-time optimization adds 18 to 20 % of revenue when the production forecast is imperfect and 7.1 % when production is known in advance. The MPC and the forecast therefore act on the same error.

7. References

Jérémie Bottieau, Louis Hubert, Zacharie De Grève, François Vallée, and Jean-François Toubeau. Very-short-term probabilistic forecasting for a risk-aware participation in the single price imbalance settlement. *IEEE Transactions on Power Systems*, 35(2):1218–1230, 2020.

Thomas Carrière. *Towards seamless value-oriented forecasting and data-driven market valorisation of photovoltaic production*. PhD thesis, Université Paris sciences et lettres, MINES ParisTech, 2020. NNT 2020UPSLM019, HAL tel-02988233.

Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, volume 30, pages 3146–3154, 2017.

Jesus Lago, Grzegorz Marcjasz, Bart De Schutter, and Rafał Weron. Forecasting day-ahead electricity prices: a review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293:116983, 2021.

Pierre Pinson, Christophe Chevallier, and George N. Kariniotakis. Trading wind generation from short-term probabilistic forecasts of wind power. *IEEE Transactions on Power Systems*, 22(3):1148–1156, 2007.

Rafał Weron. Electricity price forecasting: a review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4):1030–1081, 2014.