24th International Conference on Renewable Energies and Power Quality (ICREPQ'26)
Santander (Spain), 10th to 12th June 2026
ISBN-13: 978-84-09-82044-3

Impact of Retail Tariff Choice and Virtual Wallet on Forecast-Based PV Sharing in Residential Energy Communities

J. Serrano González[1], J.T. Villalonga Palou[1], J.M. Riquelme Santos[1], J.M Roldán Fernández[1]

[1] Department of Electrical Engineering
University of Seville, Avenida de los Descubrimientos S/N, 41092, Seville, Spain, javierserrano@us.es

Abstract. Collective self-consumption schemes typically require dynamic sharing coefficients—hourly allocation ratios per customer—to be submitted during the month preceding settlement, making reliable month-ahead, hour-by-hour demand forecasts essential. This paper presents a practical framework that generates these forecasts from historical consumption data and derives compliant allocation coefficients, while explicitly assessing how the economic value of forecasting depends on the retail tariff and compensation rules. Several lightweight persistence-based predictors are used to capture recent individual load behavior, and their outputs are combined through an ensemble that learns data-driven weights from past errors. The framework is evaluated on real hourly data from 18 residential customers under three representative tariff settings: PVPC, a market-based contract without virtual wallet, and a market-based contract with virtual wallet. Results show that the ensemble consistently improves upon individual predictors and achieves annual bills close to an oracle benchmark, while revealing that tariff design—especially whether export credits can offset fixed monthly charges—materially affects both total costs and the sensitivity of outcomes to PV size and export price.

Key words. collective self-consumption, demand forecasting, electricity tariffs, ensemble learning.

1. Introduction

Distributed renewables are accelerating the transformation of power systems, and collective self-consumption (CSC) has become a prominent model by enabling multiple consumers to share the benefits of a common photovoltaic (PV) plant. In Spain, the regulatory framework (Royal Decree 244/2019 and later updates) requires PV sharing coefficients to be submitted in advance to the distribution system operator; these coefficients can be fixed or time-varying, but in the dynamic option they may change hourly while still being notified ex ante.

This makes forecasting central to operationally feasible dynamic allocation. In this work, dynamic coefficients are derived from hourly, ex-ante demand forecasts issued for four-month blocks, covering all hours of the upcoming settlement period. Forecast accuracy therefore directly affects both the community's total bill and the distribution of costs and benefits across participants, with the magnitude of the impact depending on the settlement rules and retail tariff design (e.g., export credit treatment, virtual wallet carry-over, and whether credits can offset fixed power charges).

Research on collective self-consumption has grown rapidly over the last few years, driven by both new regulatory frameworks and the emergence of energy communities as practical deployment models. Initial studies largely highlighted the role of shared PV in increasing the local uptake of renewable generation and in delivering broader societal gains beyond bill reduction. In this line, Luthander et al. [1] provide a detailed overview of PV self-consumption modelling approaches and discuss why shared configurations are relevant to improve local integration of distributed generation.

As the field matured, a substantial body of work began to examine how shared energy should be allocated among participants. Minuto and Lanzini [2] review and compare a range of energy-sharing mechanisms across ownership structures and demand patterns, showing that seemingly simple design choices—such as allocating proportionally to demand, using fixed shares, or distributing benefits—can lead to markedly different outcomes at both individual and community level. Reis et al. [3] further emphasize that the "best" allocation depends on the community objective: strategies aimed at minimizing total cost are not necessarily aligned with maximizing self-consumption, and the resulting welfare distribution across members may differ even when overall performance improves. Related contributions such as Li and Okur [4] illustrate that allocation and investment decisions are interdependent: financing structures (e.g., third-party ownership versus member-funded PV) can shift feasibility conditions and reshape how benefits accrue. In parallel, Kortetmäki et al. [5] show that metering and crediting arrangements materially affect self-consumption rates and revenues, suggesting that allocation rules cannot be assessed in isolation from the measurement protocol and the valuation of exported energy.

A complementary stream focuses specifically on the Spanish context under RD 244/2019, using empirical data to quantify the consequences of alternative allocation choices. Llera-Sastresa et al. [6] report that both the degree of production–consumption alignment and the selected sharing scheme can strongly influence project viability and the distribution of costs within the group. Villalonga Palou et al. [7] study the estimation and optimization of static distribution coefficients using hourly consumption and realistic inputs (prices and irradiation), and find consistent economic gains relative to default coefficient settings.

Prior work has studied CSC allocation rules, but there is limited evidence on how the value of forecast-based dynamic coefficients depends on retail tariff choice and settlement rules (e.g., whether export credits can offset fixed charges and whether unused credits can be carried forward via a virtual wallet). This paper fills that gap with a fully implementable forecasting-to-allocation pipeline that generates hourly, ex-ante coefficients from four-month-ahead demand forecasts computed in rolling 4-month blocks, aligned with current Spanish submission requirements, and evaluates it under PVPC and two free-market contracts (with and without virtual wallet) against individual predictors, an error-weighted ensemble, and an oracle benchmark.

2. Methodology overview

The proposed CSC methodology follows a practical workflow using data typically available to community managers. Historical hourly demand is used to train lightweight forecasting models and generate forward predictions for each participant over the required submission horizon. These forecasts are then converted into a complete hourly schedule of sharing coefficients that satisfies regulatory constraints and can be directly submitted to the operator. The resulting pre-submitted coefficients are subsequently used to simulate settlement under alternative tariff and compensation rules. The analysis considers PVPC, Spain's regulated retail tariff with hourly-varying energy prices linked to wholesale signals and time-dependent regulated network charges, and benchmarks it against a representative free-market tariff (FMT) with fixed annual import and export prices (adopted as a transparent reference relative to more complex time-of-use offers). For the free-market setting, two settlement variants are examined: FMT-noVW, where unused export credits expire within the billing month, and FMT-VW, where remaining credits are carried forward via a virtual wallet (VW).

In a CSC setting, the hourly PV generation G_t produced by the shared installation must be distributed among the N members of the community. This allocation is defined through a set of sharing coefficients $\alpha_{i,t}$, which represent the fraction of G_t assigned to participant i at hour t. The resulting allocation $G_{i,t}$ is therefore

$$G_{i,t} = \alpha_{i,t} \cdot G_t \tag{1}$$

subject to the feasibility constraint

$$\sum_{i=1}^{N} \alpha_{i,t} = 1 \qquad \forall t \tag{2}$$

If fixed coefficients are adopted, $\alpha_{i,t}$ remains constant over time (typically for the whole month). Dynamic coefficients, instead, allow $\alpha_{i,t}$ to vary with t, enabling the allocation to better track heterogeneous demand profiles. This added flexibility can improve the perceived fairness of the scheme and, in many cases, reduce systematic mismatches between assigned PV energy and actual consumption.

In Spain, dynamic coefficients must be communicated to the distribution system operator in advance for all hours of the forthcoming settlement period. As a consequence, $\alpha_{i,t}$ cannot be computed from realized demand, but must be derived ex ante from expected consumption. In this work, coefficients are obtained from month-ahead, hour-by-hour demand forecasts computed in four-month blocks (fully compatible with the current Spanish submission requirement), and then used unchanged during settlement.

The procedure is summarized as follows:

1. Forecast individual demand. For each participant i, produce an hourly forecast $\hat{C}_{i,t}$ for all hours t in the upcoming period.
2. Compute the aggregated demand forecast.

$$\hat{C}_{tot,t} = \sum_{i=1}^{N} \hat{C}_{i,t} \tag{3}$$

3. Build dynamic sharing coefficients by normalization:

$$\alpha_{i,t} = \frac{\hat{C}_{i,t}}{\hat{C}_{tot,t}} \tag{4}$$

4. Allocate PV generation using the pre-submitted schedule. During the subsequent settlement period, the allocation is performed with the submitted $\alpha_{i,t}$ values, i.e., $G_{i,t} = \alpha_{i,t} G_t$, and billing is computed accordingly.

This normalization-based rule allocates PV output proportionally to each member's expected demand at every hour. In doing so, it provides a transparent and directly implementable mechanism to translate forecasts into dynamic coefficients, with the aim of limiting allocation imbalances and improving the distribution of costs and benefits within the community.

3. Forecasting approaches

In this study, several lightweight persistence-style predictors are implemented to capture different aspects of recent consumption behavior. Their month-ahead outputs—generated in rolling four-month blocks—are combined through an ensemble (ENS) whose hour-of-week weights are learned from recent forecast errors using a strictly backward validation window, thus avoiding any look-ahead bias. The aim of the ensemble is not to enforce a specific theoretical property, but to reduce forecasting risk relative to any individual predictor by adapting the combination weights to recent performance. The resulting forecasts are then translated into ex-ante dynamic sharing coefficients, which are submitted ahead of settlement and used to assess how improved demand anticipation affects the economic performance of shared PV generation at community level under different tariff schemes.

A. Seasonal profile model (SP)

SP is a non-parametric, persistence-based predictor that exploits the strong 168-hour (hour-of-week) seasonality of residential demand. It builds a 168-point weekly template from past observations (with mild winsorization to limit outliers), rescales the template to match the recent consumption level (ratio of recent mean to long-term mean), and then repeats the adjusted weekly profile forward to obtain the month-ahead forecast (clipping negative values to zero). The method is lightweight, interpretable, and performs best when demand is dominated by stable weekly cycles, while atypical events (e.g., holidays or abrupt changes) are not modeled explicitly.

B. Previous-month hour-of-day model (PM)

The PM model is a persistence-based predictor that assumes the most recent month contains the most relevant information about near-future consumption patterns. Instead of relying on weekly repetition, PM focuses on the hour-of-day structure observed during the last complete month available before the forecast is issued. For each hour of the day, a representative demand value is obtained from the observations recorded at that same hour in the previous month, after applying light filtering to reduce the influence of atypical days or missing samples. This 24-hour daily profile is then repeated forward to produce the full month-ahead hourly forecast, starting from the next hour after the last available measurement. Negative values are truncated to zero. Because it relies only on very recent data, PM reacts quickly to gradual behavioral changes, although it does not explicitly capture longer-term seasonal effects.

C. Same calendar month in previous years model (PY)

The PY model captures seasonal recurrence by assuming that consumption patterns observed in the same calendar month of previous years are indicative of the upcoming month. Historical data from all past occurrences of that calendar month are aligned by hour-of-day, and for each hour a representative value is computed to form a daily seasonal profile. This 24-hour profile is then repeated to generate the month-ahead hourly forecast, beginning immediately after the last observed sample, with negative values set to zero. PY effectively reflects recurring seasonal influences—such as climatic or habitual patterns—that short-term persistence models may overlook, although it may be less responsive to recent shifts in consumption behavior.

D. Rolling multi-month hour-of-day model (ROLL)

The ROLL model is a persistence-based predictor designed to balance short-term adaptation with medium-term stability. Instead of relying on a single recent month or on seasonal recurrence from previous years, ROLL builds its forecast from the aggregated behavior observed over a rolling window covering several recent months.

All historical observations within this rolling window are grouped by hour-of-day, and for each of the 24 hours a representative demand value is computed after lightly filtering atypical samples. This produces a daily profile that reflects the average behavior over multiple recent months, smoothing out temporary anomalies while still remaining sensitive to gradual changes in consumption habits. The resulting 24-hour profile is then repeated forward from the next available hour to generate the complete month-ahead forecast, with negative values clipped to zero.

By combining information from multiple recent months, ROLL provides a compromise between the rapid adaptability of PM and the long-term seasonal view of PY, often yielding stable and robust forecasts when consumption patterns evolve slowly over time.

E. K-Nearest Neighbors (KNN) Forecasting Method

The KNN forecaster is a data-driven, non-parametric approach that predicts each hour of the target month by searching for *historical situations that look similar* to the most recent consumption pattern of each customer. For every forecasted hour, a feature vector is built from recent load information (e.g., lagged hourly values and simple calendar descriptors such as hour-of-day and day-of-week), and the algorithm retrieves the $\boldsymbol{K}$nearest historical observations according to a distance metric defined on those features. The prediction is then obtained as an average (or distance-weighted average) of the corresponding realized loads of those neighbors, which allows the model to adapt to recurring—but not strictly periodic—behavioral regimes. Because it relies only on past measured demand and local similarity rather than an explicit parametric structure, KNN can capture nonlinear patterns and atypical weeks better than simple persistence in some cases, although its performance depends on the availability of comparable past examples and it can be sensitive to the choice of features, distance metric, and $\boldsymbol{K}$.

F. Ensemble forecasting approach (ENS)

Each of the individual predictors captures a different aspect of consumption persistence. SP exploits weekly regularity, PM reacts to very recent behavior, PY reflects seasonal recurrence, and ROLL smooths patterns over several recent months. This study adopts an adaptive ensemble scheme, building on our earlier work [8], to leverage the complementary strengths of the base predictors, limit reliance on any single model, and update the combination as recent conditions change.

For every issuance date and for each user, the ensemble operates at the level of hour-of-week strata $\boldsymbol{h} \in \{\mathbf{1}, \dots, \mathbf{168}\}$. This allows the combination to account for systematic differences in model performance across hours of the week. All weights are estimated using only data available up to the end of the previous month, ensuring a strictly backward-looking validation window with no look-ahead bias.

Let $\hat{\boldsymbol{C}}_{\boldsymbol{m},\boldsymbol{h},\boldsymbol{t}}^{(\boldsymbol{j})}$ denote the forecast of *j*-th model for a future hour $\boldsymbol{t}$ belonging to stratum $\boldsymbol{h}$ in a issuance month $\boldsymbol{m}$. The ensemble forecast for that hour is defined as a weighted combination of the base predictors:

$$\hat{C}_{m,h,t}^{ENS} = \sum_{j=1}^{J} w_{m,h}^{(j)} \hat{C}_{m,h,t}^{(j)} \tag{5}$$

where $\boldsymbol{w}_{\boldsymbol{m},\boldsymbol{h}}^{(\boldsymbol{j})}$are the weights associated with model $\boldsymbol{j}$at hour-of-week stratum $\boldsymbol{h}$. To estimate these weights, past forecast errors are computed within a rolling calibration window that ends at the issuance cut-off (the last hour of the previous month). For each model $\boldsymbol{j}$and stratum $\boldsymbol{h}$, the historical errors are:

$$e_{h,t}^{(j)} = C_t - \hat{C}_t^{(j)} \quad (6)$$

where $\boldsymbol{C_t}$ is the observed demand. From these errors, the empirical variance of each model in that stratum is estimated as:

$$\sigma_h^{2(j)} = \frac{1}{N_h} \sum_{t \in h} \left(e_{h,t}^{(j)} \right)^2 \quad (7)$$

with $\boldsymbol{N_h}$ being the number of samples in stratum $\boldsymbol{h}$.

The ensemble weights are chosen to minimize the variance of the combined forecast subject to a sum-to-one constraint:

$$\min_{w_h^{(j)}} \sum_{j=1}^{J} \left(w_h^{(j)} \right)^2 \sigma_h^{2(j)} \quad \text{s.t.} \quad \sum_{j=1}^{J} w_h^{(j)} = 1 \quad (8)$$

Solving this constrained problem yields inverse-variance weights normalized to one:

$$w_h^{(j)} = \frac{\left(\sigma_h^{2(j)} \right)^{-1}}{\sum_{k=1}^{J} \left(\sigma_h^{2(k)} \right)^{-1}} \quad (9)$$

Finally, these hour-of-week weights are applied to the base month-ahead forecasts to obtain the ensemble prediction used for computing the dynamic sharing coefficients:

$$\hat{C}_{m,t}^{ENS} = \sum_{j=1}^{J} w_{m,h(t)}^{(j)} \hat{C}_{m,t}^{(j)} \quad (10)$$

where $\boldsymbol{h(t)}$ denotes the hour-of-week stratum to which hour $\boldsymbol{t}$ belongs.

By adapting the combination weights to recent performance at each hour of the week, the ENS forecast preserves temporal structure and provides a more robust basis for defining the ex-ante dynamic sharing coefficients required in CSC settlement.

4. Test cases

This section evaluates the proposed forecasting-and-allocation pipeline on a real collective self-consumption (CSC) community of 18 residential customers with hourly metered demand, using the 2021 settlement year as the evaluation period for both forecasting accuracy and billing outcomes. Forecasts are generated in a strictly backward-looking manner (using only data available up to each issuance time) and are produced month-ahead at hourly resolution through a rolling four-month block workflow, enabling the preparation of a complete schedule of hourly sharing coefficients for the required submission horizon without future information.

The economic assessment considers three tariff settings representative of common Spanish settlement designs. The first is PVPC, the regulated retail tariff, where the purchase price varies hourly with wholesale market signals and regulated components, including time-dependent network tolls and charges. Two free-market reference tariffs are also analysed assuming constant import and export prices (acknowledging that time-of-use commercial offers exist): FMT-noVW, where export credits are limited to the current billing month, and FMT-VW, where unused credits can be carried forward through a virtual wallet. Consistent with the implemented settlement logic, PVPC export credits cannot offset fixed monthly charges, while the wallet-enabled design can enable inter-month compensation.

The shared generation asset is modelled as a rooftop PV plant in Seville (Spain) with a base rated power of 35 kW, using an hourly production profile scaled from a per-kW reference series. Hourly energy balances and monthly settlements combine (i) realized demand, (ii) PV generation, (iii) tariff-specific import/export prices, and (iv) regulated fixed charges where applicable. For the market-based cases, reference prices are set to 12 c€/kWh (imports) and 7 c€/kWh (exports), matching the script configuration.

The case study quantifies how forecast-driven dynamic sharing affects settlement outcomes under different tariff and compensation rules. Several lightweight persistence-style predictors are considered for month-ahead hourly forecasting. Their outputs are combined through the proposed ENS approach, which learns hour-of-week–dependent weights from recent forecast errors using a strictly backward validation window. Forecasts are converted into ex-ante hourly sharing coefficients via normalization and used to allocate PV generation during settlement. Results compare individual predictors, ENS, and an oracle benchmark (based on realized demand) in terms of monthly billing and annual total costs, and include sensitivity analyses on PV size and export

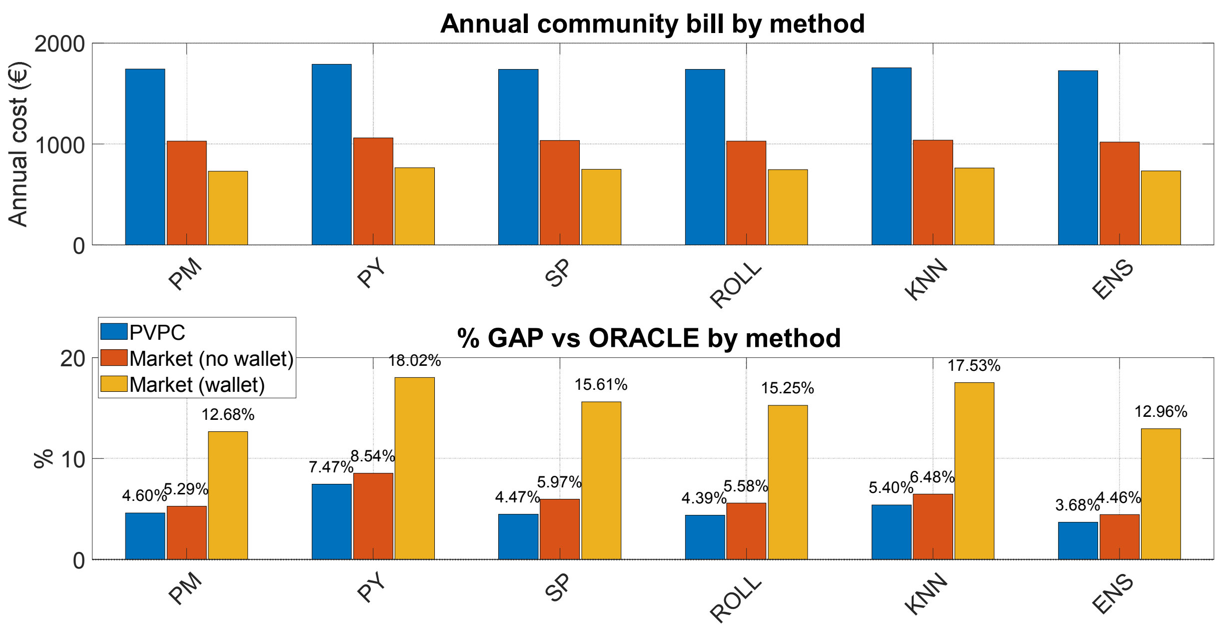

Fig.1. Per-tariff comparison (predictors vs ENS vs Oracle): Annual Bill by Forecast Method (vs Oracle). % Gap to Oracle Shown Below

remuneration.

Figure 1 compares the annual community bill obtained when dynamic sharing coefficients are computed from each individual predictor and from the proposed ensemble (ENS), for each of the three tariff settings considered. Results are benchmarked against an oracle reference, which represents an idealized case where each participant's realized consumption is assumed to be known in advance and therefore coefficients can be computed without forecasting uncertainty. Overall, the proposed ensemble consistently delivers the best performance across tariffs, improving upon all individual predictors. In PVPC, ENS achieves the smallest deviation from the oracle, with a 3.68% gap. Under the free-market retail tariff without a Virtual Wallet, the ENS–oracle difference is 4.46%. When a Virtual Wallet is enabled, the ENS gap increases to 12.96%, yet it remains lower than the gaps produced by any single forecasting method.

Table 1 reports the community's total monthly bills and the annual total for 2021 under the three tariff options—PVPC, FMT-noVW, and FMT-VW—computed with the proposed ENS approach; the corresponding oracle values are shown in parentheses as an ideal reference where participants' consumptions are assumed to be known in advance. In annual terms, ENS yields €1,726.9 for PVPC (oracle: €1,665.5), compared with €1,019.3 for FMT-noVW (oracle: €975.7) and €732.5 for FMT-VW (oracle: €648.5). This confirms that PVPC is the least favorable option economically for this community, while both free-market settings substantially reduce total costs. The table also illustrates the added value of the virtual wallet design: by allowing unused export credits to be carried over to subsequent months, FMT-VW avoids losing credits accumulated during high-radiation periods and achieves the lowest annual bill—e.g., ENS drops from €153.1 in November (FMT-noVW) to €0 under FMT-VW, and from €208.9 to €163.8 in December, reflecting the inter-month compensation enabled by the wallet mechanism.

Table I. – Annual electricity bill (€) depending on tariff choice obtained by the proposed ENS method (theoretical maximum obtained by oracle approach shown in brackets)

Month	PVPC	FMT-noVW	FMT-VW
Jan	341.0 (332.1)	271.7 (263.2)	271.7 (263.2)
Feb	202.3 (197.9)	148.2 (143.9)	148.2 (143.9)
Mar	189.2 (178.7)	131.5 (122.6)	131.5 (122.6)
Apr	105.0 (94.8)	17.4 (10.2)	17.4 (10.2)
May	62.5 (53.7)	0 (0)	0 (0)
Jun	47.3 (47.3)	0 (0)	0 (0)
Jul	47.5 (47.5)	0 (0)	0 (0)
Aug	47.5 (47.5)	0 (0)	0 (0)
Sep	47.3 (47.3)	0 (0)	0 (0)
Oct	143.8 (138.3)	88.5 (84.2)	0 (0)
Nov	221.8 (215.7)	153.1 (148.4)	0 (0)
Dec	271.6 (264.5)	208.9 (203.3)	163.8 (108.7)
TOTAL	**1726.9 (1665.5)**	**1019.3 (975.7)**	**732.5 (648.5)**

Figure 2 shows the annual electricity bill of each of the 18 customers under the three tariff options considered (PVPC, FMT-noVW, and FMT-VW) using the proposed ensemble-based allocation (top panel). The bottom panel reports, for each customer and tariff, the percentage gap with respect to the oracle benchmark, i.e., the ideal case where customers' actual demand is assumed to be known in advance when computing the sharing coefficients. Overall, the figure highlights both the strong customer-to-customer variability in annual costs and how the tariff design and compensation rules shape the sensitivity to forecast-driven allocation: PVPC typically leads to higher bills, while the free-market settings reduce costs, and the virtual wallet option further improves outcomes by allowing surplus-export credits accrued during high PV

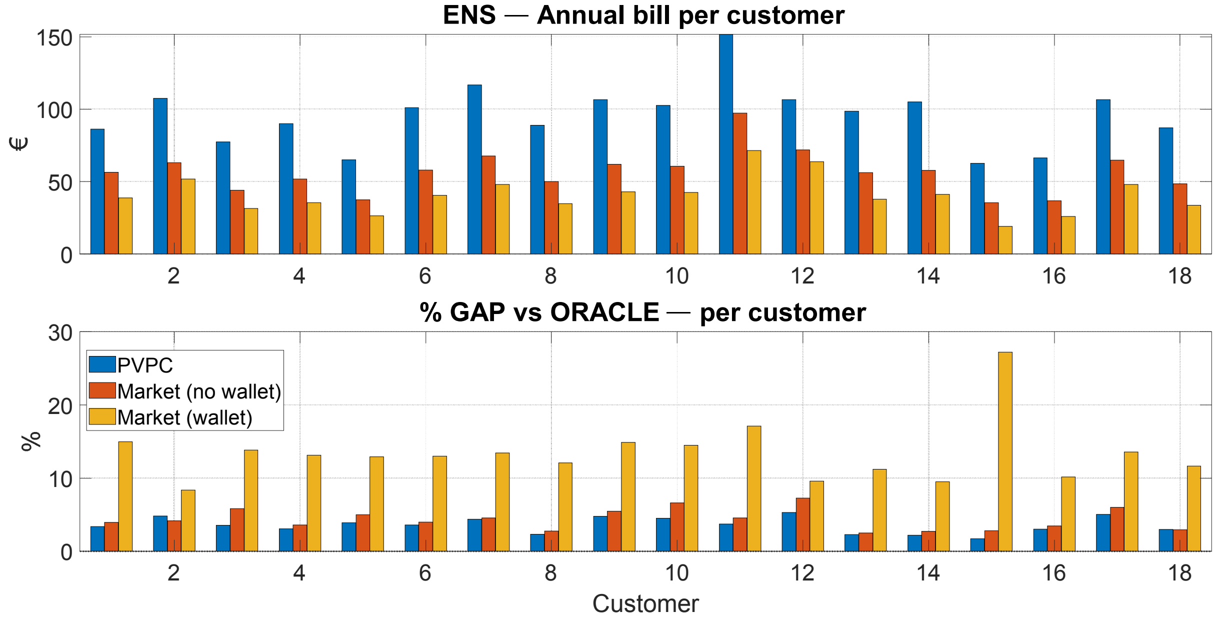

Fig.2. Annual bill per customer (€) across tariff options (top) and percentage gap vs. oracle benchmark (bottom)

generation periods to be carried over and used in subsequent months.

A. Sensitivity analysis

This section performs a sensitivity analysis to assess the robustness of the results by varying the PV system size. Specifically, the rated PV power is swept from 75% to 125% of the baseline value adopted in the previous section.

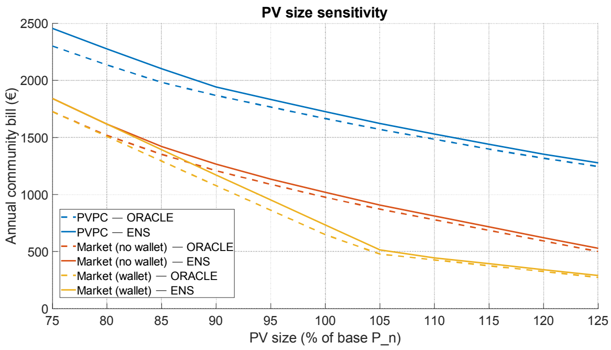

Fig.3. PV size sensitivity (75–125% of baseline): annual community bill (ENS vs oracle).

Fig. 3 reports a PV size sensitivity analysis in which the rated PV power is scaled from 75% to 125% of the 35 kW baseline. For each scaling factor, the annual community bill is recomputed under the three tariff settings, comparing the proposed ensemble-based allocation (ENS) against the oracle benchmark (perfect foresight of demand). The curves show that increasing PV capacity reduces the total annual cost, while the magnitude of the savings and the remaining gap to the oracle depend on the tariff design and its compensation rules. In particular, the virtual wallet has a negligible impact below 80% of 35 kW (i.e., 28 kW), whereas for larger PV sizes it becomes increasingly relevant by allowing surplus credits to be carried over and thus monetized more effectively across months.

5. Conclusion

The paper concludes that forecast-driven, dynamic PV sharing can be implemented in a fully realistic way under current Spanish CSC requirements by generating month-ahead hourly sharing coefficients from rolling four-month hourly demand forecasts, and that an adaptive error-weighted ensemble (ENS) consistently outperforms the individual persistence-style predictors considered (SP, PM, PY, ROLL, and KNN), yielding annual community bills that remain close to an oracle benchmark based on perfect knowledge of realized demand. Beyond this methodological contribution, the results emphasize that the economic value of forecasting is highly contingent on retail tariff design and settlement rules (often to the extent that these design choices condition, or even dominate, the improvements attributable to better predictions). In particular, PVPC is systematically the least favorable option for the studied community, while free-market tariffs reduce total costs under the same physical demand and PV generation, indicating that identical allocation/forecasting strategies can lead to materially different outcomes depending on how imports, exports, and regulated components are priced and settled. The analysis further shows that enabling a virtual wallet strengthens the benefits of dynamic sharing by allowing surplus-export credits to be carried across months, smoothing the mismatch between periods of high PV output and periods of high consumption. This wallet effect is negligible at low PV penetration, specifically for capacities below roughly 28 kW (80% of the 35 kW baseline), because monthly surplus is limited and credits are largely exhausted within the same billing period; however, as PV size increases and surplus generation becomes more frequent, inter-month carryover becomes progressively more valuable and translates into a visibly lower annual bill.

The findings suggest two key takeaways: first, robust forecasting-to-allocation pipelines (and especially adaptive ensembling) do improve performance and help approach the oracle upper bound; second, the largest swings in billing outcomes are driven by settlement mechanics—how export credits are capped, whether credits can offset fixed charges, and whether unused credits can be transferred to future months—which ultimately govern how much of the PV value can be monetized and how sensitive results are to PV sizing and export remuneration.

Acknowledgement

This work has been supported by MICIU/AEI/ 10.13039/501100011033 and by ERDF/EU under grant PID2021127550OA–I00.

References

[1] R. Luthander, J. Widén, D. Nilsson, and J. Palm, “Photovoltaic self-consumption in buildings: A review,” *Appl. Energy*, vol. 142, pp. 80–94, Mar. 2015, doi: 10.1016/j.apenergy.2014.12.028.

[2] F. D. Minuto and A. Lanzini, “Energy-sharing mechanisms for energy community members under different asset ownership schemes and user demand profiles,” *Renew. Sustain. Energy Rev.*, vol. 168, p. 112859, Oct. 2022, doi: 10.1016/j.rser.2022.112859.

[3] I. F. G. Reis, I. Gonçalves, M. A. R. Lopes, and C. H. Antunes, “Collective self-consumption in multi-tenancy buildings–To what extent do consumers’ goals influence the energy system’s performance?,” *Sustain. Cities Soc.*, vol. 80, p. 103688, May 2022, doi: 10.1016/j.scs.2022.103688.

[4] N. Li and Ö. Okur, “Economic analysis of energy communities: Investment options and cost allocation,” *Appl. Energy*, vol. 336, p. 120706, Apr. 2023, doi: 10.1016/j.apenergy.2023.120706.

[5] A. Kortetmäki, J. Ylipaino, J. Koskela, K. Kalliohariju, and P. Järventausta, “The impact of metering methods on collective self-consumption: Insights from multi-dwelling buildings in Finland,” *Energy Build.*, vol. 319, p. 114511, Sep. 2024, doi: 10.1016/j.enbuild.2024.114511.

[6] E. Llera-Sastresa, J. Á. Gimeno, J. L. Osorio-Tejada, and P. Portillo-Tarragona, “Effect of Sharing Schemes on the Collective Energy Self-Consumption Feasibility,” *Energies*, vol. 16, no. 18, p. 6564, Jan. 2023, doi: 10.3390/en16186564.

[7] J. T. Villalonga Palou, J. Serrano González, and J. M. Riquelme Santos, “Estimation of Energy Distribution Coefficients in Collective Self-Consumption Using Meta-Heuristic Optimization Techniques,” *Sustainability*, vol. 16, no. 7, p. 2741, Jan. 2024, doi: 10.3390/su16072741.

[8] J. T. Villalonga Palou, J. Serrano González, J. M. Riquelme Santos, and J. M. Roldán Fernández, “A novel weight-based ensemble method for emerging energy players: an application to electric vehicle load prediction,” *Energy AI*, vol. 20, p. 100510, May 2025, doi: 10.1016/j.egyai.2025.100510.