Funded by the European Union

Horizon Europe

EUROPEAN COMMISSION

European Climate, Infrastructure and Environment Executive Agency (CINEA)

Grant agreement no. 101160684

U2DEMO

Use of open-source P2P energy sharing platforms for energy Democratization

Deliverable D 2.2
Decision support methods for active consumers and ECs

Document Details

Due date	31-10-2025
Actual delivery date	31-10-2025
Lead Contractor	R&D Nester
Version	1.0
Prepared by	Inoussa Laouali (R&D Nester), Gonçalo Glória (R&D Nester), Larissa da Silva Montefusco (INESC ID), Diogo Pinto Ferreira (INESC ID), Hugo Morais (INESC ID), Jaisiel Santana (EIFER), Maxime Gatineau (EIFER), Hubert Kriebitzsch (EIFER), Geert Deconinck (KU Leuven)
Teams involved	Thomas Heylen and Steven Laurijssen (Klimaan)
Reviewed by	Humberto Queiroz (EDP NEW) and Ivan Pasarin (EWF) Hossein Fani, Wicak Ananduta and Anibal Sanjab (VITO) Glenn Reynders (Scientific coordinator review)
Dissemination Level	Public

Project Contractual Details

Project Title	Use of open-source P2P energy sharing platforms for energy democratization
Project Acronym	U2Demo
Grant Agreement No.	101160684
Project Start Date	01-09-2024
Project End Date	29-02-2028

Duration	42 months

Document History

Version	Date	Contributor(s)	Description
0.1	06.06.2025	Gonçalo Glória et al.	Initial draft TOC
0.11	03.08.2025	Gonçalo Glória et al.	Chapter 6 - draft
0.2	27.08.2025	Gonçalo Glória et al.	Chapter 5 – draft
0.3	15.09.2025	Gonçalo Glória et al.	Chapter 3 – draft
0.4	25.09.2025	Gonçalo Glória et al.	Chapter 2– draft
0.5	07.10.2025	Gonçalo Glória et al.	Chapter 4 – draft
0.6	10.10.2025	Gonçalo Glória et al.	Final version for internal review
0.7	16.10.2025	Humberto Queiroz	First review
0.71	22.10.2025	Ivan Pasarin	Second review
0.72	24.10.2025	Hossein Fani, Wicak Ananduta and Anibal Sanjab	WP leader review
0.8	27.10.2025	Gonçalo Glória	Final version for scientific coordinator review
0.9	29.10.2025	Glenn Reynders	Scientific coordinator review
1.0	31.10.2025	Coordination team	Final version

Disclaimer

This document has been produced in the context of the U2Demo[1] project. Views and opinions expressed in this document are however those of the authors only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the grating authority can be held responsible for them.

Acknowledgment

This document is a deliverable of U2Demo project. U2Demo has received funding from the European Union´s Horizon Europe programme under grant agreement no. 101160684.

Funded by the European Union

[1] https://u2demo.eu/

Executive Summary

Deliverable D2.2 presents the development and proof of concept of four complementary methods designed to support active consumers and energy communities (ECs) in defining and optimising their operational strategies. These methods, Machine-Learning based Forecasting, Data Analytics, Decision Support, and Distributed Energy Resource (DER) Valorisation, form a coherent framework enabling data-driven participation in local energy sharing and peer-to-peer (P2P), market interaction, and flexibility services. The deliverable, developed under Task 2.2 of the U2Demo project, aims to demonstrate the technical feasibility and complementarity of the four methods, validating their capability to support informed, optimised, and customised decision-making processes for consumers and ECs.

The Machine-Learning based Forecasting methods, developed by EIFER, provide predictive tools for electricity demand, renewable generation (PV and wind), electric vehicle charging needs, and market price signals. By incorporating exogenous variables such as weather data, time features, and historical patterns, these models deliver accurate short- and medium-term forecasts tailored to different levels of granularity (individual, EC, or regional). The forecasting results constitute a key input for the remaining methods, as they enable proactive planning and flexibility assessment.

The Data Analytics methods, developed by KU Leuven, evaluate consumer and community behaviour across technical, economic, social, and environmental dimensions. They quantify how energy strategies influence self-consumption, self-sufficiency, cost savings, fairness, and carbon reduction. The analytics framework combines descriptive, predictive, and prescriptive approaches to create a unified evaluation structure. In this deliverable, the methods have been applied using reference datasets and synthetic energy community profiles to establish baselines and assess the expected impact of different operational strategies. Importantly, Data Analytics also provide validation feedback for forecasting models and serve as a foundation for defining meaningful KPIs used in subsequent optimisation tasks.

The Decision Support methods, developed by INESC ID, offer a structured optimisation-based framework for energy community management. Implemented in the PyECOM platform, these methods evaluate operational strategies under multiple objectives, such as reducing operational cost, electricity cost, or grid import, while improving comfort, battery longevity, and environmental impact. Multi-objective optimisation techniques are employed to assess trade-offs between conflicting objectives and identify balanced operational strategies. Results demonstrate the capacity of the decision support system to align community-wide resource coordination with diverse member preferences, providing actionable recommendations to EC managers and supporting participation in P2P trading and flexibility markets.

The DER Valorisation methods, developed by R&D Nester, address peer-level optimisation and the quantification of flexibility potential from individual distributed resources such as PV systems, stationary batteries, Electrical Vehicles (EVs), and Heating, Ventilation and Air Conditioning (HVAC) loads. The approach is formulated as a multi-period optimisation problem that determines optimal DER scheduling while respecting technical constraints and comfort conditions. The tool evaluates flexibility margins and economic value of DER operation, aiming to improve self-consumption, minimise curtailment, and create added value through potential participation in market or P2P schemes. The results demonstrate that coordinated operation can significantly reduce grid dependency, increase the use of renewable energy, and add economic return by DER participation.

Although each method has been implemented and validated independently, their design follows an integrated methodological framework that ensures data and functional

interoperability. Forecasting provides the predictive signals required by the other methods, Data Analytics validates model behaviour and defines performance indicators, and both Decision Support and DER Valorisation use these inputs to generate optimised operational strategies. This framework establishes the logical flow of information and retroactive feedback loops necessary for future integration within the U2Demo platform.

Deliverable D2.2 thus represents the proof of concept of the four methods, serving as a foundation for subsequent project activities. In Task 4.2, the focus will shift towards increasing the Technology Readiness Level (TRL) of these methods through their implementation and integration in the U2Demo operational platform. The outputs defined in D2.2 will also serve as inputs to Task 2.3, where energy sharing and P2P models are being developed, and to Task 2.4, which addresses market and flexibility participation mechanisms, including system services for Transmission and Distribution System Operators (TSOs and DSOs).

In summary, this deliverable establishes a solid methodological basis for empowering energy communities and active consumers through data-driven decision-making. The four methods complement each other by linking forecasting accuracy, behavioural insight, optimisation capability, and flexibility valorisation. Together, they form an integrated set of tools that will support the transition towards interoperable, transparent, and participatory energy systems under the U2Demo framework.

Table of Contents

List of Figures

List of Tables

Keywords, Acronym

ABAC	Attribute-based access control
AES	Advanced Encryption Standard
AIME	Agile Interaction Model based ontology development Methodology
API	Application Programming Interface
AR	Autoregressive
BACS	Building Automation and Control Systems
BEM	Block Element Modifier
BESS	Battery Energy Storage System
BIM	Building Information Modelling
BO	Business Object
BUC	Business Use Case
CIM	Common Information Model
CPU	Central Processing Unit
CS	Cost savings
CSV	Comma-separated values
DA	Day-Ahead
DER	Distributed Energy Resource
DID	Decentralized Identifier
DL	Deep Learning
DR	Demand Response
DS	Data Space
DSO	Distribution System Operator
EC	Energy Community
EMS	Energy Management Systems
ETS	Error, Trend and Seasonality
EU	European Union
EV	Electric Vehicle
EVSE	Electric Vehicle Supply Equipment
EW	Energy Web
FAIR	Findable, Accessible, Interoperable, and Reusable
GBM	Gradient Boosting Machine
GDPR	General Data Protection Regulation
GEMS	Grid Energy Management System
HBES	Home and Building Electronic Systems
HEMS	Home Energy Management System
HP	Heat Pump
HTTP	Hypertext Transfer Protocol
IEC	International Electrotechnical Commission
IoT	Internet of Things
ISO	International Organization for Standardization
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MFA	Multi-factor authentication
ML	Machine Learning
MQTT	Message Queuing Telemetry Transport
MPSI	Maximum potential saving index
mTLS	Mutual Transport Layer Security
NCCS	Network Code on Cybersecurity

NIS 2	Network and Information Systems 2 Directive
NISTIR	National Institute of Standards and Technology Interagency Reports
OAuth2	Open Authorization Protocol
OCPI	Open Charge Point Interface
OCPP	Open Charge Point Protocol
OPENA	Open Automated Demand Response
OWL	Web Ontology Language
P2P	Peer-to-peer
PKI	Public key infrastructure
PV	Photovoltaic
RBAC	Role-based access control
RDF/S	Resource Description Framework Schema
REST	Representational State Transfer
RM	Resource Manager
RMSE	Root Mean Squared Error
RNN	Recurrent Neural Network
SAREF	Smart Applications REFerence
SCR	Self-consumption ratio
SEAS	Smart Energy Aware Systems
sFTP	Secure File Transfer Protocol
SHACL	Shapes Constraint Language
SMPC	Secure Multi-Party Computation
SSR	Self-sufficiency ratio
TEE	Trusted Execution Environment
TFT	Temporal Fusion Transformer
TiDE	Time-series Dense Encoder
TLS	Transport Layer Security
TSO	Transmission System Operator
TSV	Tab-separated values
TTL	Terse Triple Language
UC	Use Case
URI	Uniform Resource Identifier
UTC	Coordinated Universal Time
UX	User Experience
VUF	Voltage unbalance factor
WAPE	Weighted Absolute Percentage Error
WP	Work Package
WQL	Weighted Quantile Loss
XML	eXtensible Markup Language

1 Introduction

1.1 Scope and Objectives

Task 2.2 is dedicated to the development and assessment of methods designed to support consumers and energy communities (ECs) in the definition of their operational and participation strategies. The focus is on enabling active involvement in energy markets, peer-to-peer (P2P) trading schemes, and flexibility services, through the application of data-driven and model-based approaches. The methods are intended to cover forecasting, decision support, behavioural analytics, and distributed energy resource (DER) valorisation, forming a coherent framework for enhanced consumer and community energy management.

The first methodology concerns **machine learning forecasting methods** (developed by EIFER), where AI techniques are applied to provide short-term and medium-term predictions relevant to ECs and prosumers. The forecasting scope includes demand, renewable generation (photovoltaic and wind), and electric vehicle charging requirements. In parallel, forecasts of market prices and flexibility needs are incorporated, with the objective of improving scheduling decisions and enabling informed participation in energy sharing, P2P and/or flexibility services.

The second methodology focuses on **data analytics methods** (developed by KU Leuven), which analyse consumer behaviour and the performance of different strategies in terms of consumption, production, and storage patterns. This analysis is carried out across technical, economic, and social dimensions, allowing the identification of strengths and weaknesses in adopted strategies. The results feed back into strategy refinement, improving both efficiency and effectiveness of consumer and community-level actions.

The third methodology addresses **decision support methods** (developed by INESC), which build on forecasting outputs and contextual information to generate recommendations for active consumers. These methods are designed to assist in the evaluation of participation strategies for P2P trading, market interaction, and the provision of flexibility services. By integrating technical constraints and economic considerations, decision support mechanisms provide structured guidance for optimising operational choices by an EC manager.

The fourth methodology covers **DER valorisation methods** (developed by R&D Nester), targeting the systematic evaluation of distributed energy resources in the context of energy sharing and P2P trading. The aim is to quantify the value of individual and aggregated DERs under different operating conditions, thereby providing peers with the necessary information to maximise resource utilisation.

Taken together, these four methodological components establish the basis for improved participation of consumers and ECs in energy sharing frameworks and P2P mechanisms. By enabling accurate forecasting, structured decision support, behavioural insights, and transparent resource valorisation, Task 2.2 directly contributes to the design of optimised and customised decision-making strategies. These strategies strengthen consumer and EC engagement in collaborative energy exchanges while ensuring that technical feasibility, strategic decision-making, facilitate efficient resource management, and increase engagement in market-based and community-driven mechanisms.

1.2 Structure

The deliverable is organized into seven main parts:

- **Section 1 – Introduction**
  Defines the scope, objectives, and structure of the deliverable, situating Task 2.2 within the broader U2Demo context. It highlights the relevance of the methods developed for supporting consumers and energy communities in P2P participation and energy sharing.
- **Section 2 – Methodological Framework**
  Describes the overall approach used to design and assess the four methods, including their interdependencies and intended outputs. It provides the rationale for selecting the methods and the criteria for evaluation.
- **Section 3 – Machine-Learning based Forecasting methods**
  Presents forecasting techniques developed by EIFER for consumption, renewable generation, EV demand, and market signals. It details the modelling approaches, data requirements, and expected contribution to flexibility and P2P decision-making.
- **Section 4 – Data Analytics methods**
  Explains the methods from KU Leuven for analysing consumer and community behaviour across consumption, production and storage. The section emphasises how technical, economic, and social insights support the refinement of strategies.
- **Section 5 – Decision support methods**
  Details the tools developed by INESC to provide actionable recommendations for EC participation in P2P trading, markets, and flexibility services.
- **Section 6 – DER valorisation methods**
  Describes the approaches proposed by R&D Nester to quantify and optimise the value of DERs in P2P and energy-sharing contexts.
- **Section 7– Conclusions**
  Summarises the main findings of Task 2.2 and the contribution of each methodology to consumer and EC empowerment. It also outlines the next steps and how results will be leveraged in subsequent U2Demo tasks.

1.3 Relationship with other deliverables

This document provides the proof of concept of the four methodologies developed in Task 2.2, serving as a foundation for subsequent activities within the project. The follow-up work in Task 4.2 will focus on increasing the Technology Readiness Level (TRL) of these methods and implementing them in the operational context of the U2Demo platform.

Furthermore, the methodologies defined in Deliverable 2.2 constitute a direct input to Task 2.3, where energy sharing and peer-to-peer (P2P) models are being developed. Beyond this, the outputs can also support the extension of energy community and active consumer participation to market-based and flexibility mechanisms, including the provision of system services to Transmission System Operators (TSOs), Distribution System Operators (DSOs), and other services like demand response programs, that is being addressed in Task 2.4.

2 Methodological Framework

The methodological framework developed in Task 2.2 establishes the logical and technical interconnection between the four methods: machine-learning based forecasting, data analytics, decision support, and DER valorisation. Rather than being conceived as isolated modules, these methodologies form a coherent process where inputs and outputs are exchanged, and retroactive feedback mechanisms ensure continuous improvement and alignment with the needs of consumers and ECs.

At the core of the framework are the forecasting methods, which provide time-series predictions of key variables such as electricity consumption, renewable generation (PV and wind), electric vehicle demand, and market prices. In addition, forecasts of flexibility needs are produced. These forecasting outputs act as primary inputs to the other methods. They are consumed by the decision support tools (DER valorisation and decision support) to build scenarios and recommend strategies, and by the data analytics layer to compare predicted versus observed behaviour, thereby validating and refining the models.

The data analytics methods perform a dual role in the framework. First, they analyse consumer and community behaviour by evaluating actual energy usage, generation, and storage strategies. This analysis provides feedback to the forecasting models, enabling the improvement of prediction accuracy through the integration of observed behavioural patterns. Second, analytics generate performance indicators that are used by the decision support methods to contextualise recommendations, ensuring that technical, economic, and social aspects are simultaneously considered.

The decision support methods operate as the centralised mechanism that translates forecasts and analytics into actionable strategies for consumers and ECs. By integrating forecast data, behavioural insights, and technical constraints, decision support tools evaluate alternative participation options in P2P trading, energy markets, and flexibility services. The results are customised recommendations that aim to maximise value while ensuring feasibility.

The DER valorisation methods provide a decentralised decision-support framework for optimising the operation of distributed resources such as PV, stationary batteries, EVs, and flexible loads. Through a multi-period optimisation process, the tool schedules DERs over a defined horizon while respecting technical and operational constraints, quantifying feasible flexibility margins, and maximising self-consumption and renewable integration. By balancing economic efficiency, technical feasibility, and user comfort, the methodology enables active consumers to unlock the operational and economic value of their DERs, creating a direct link to flexibility markets and P2P trading mechanisms.

It is important to note that, although this deliverable focuses on the proof of concept of each methodology as a standalone development, their potential integration has been considered from the outset. The design of each component was guided by the need to ensure interoperability and consistency, so that in the next phase (Task 4.2) these methodologies can be combined, scaled, and implemented within the U2Demo platform.

Overall, the methodological framework is cyclical and interdependent:

- Forecasting generates the predictive signals needed for proactive planning and decision support.
- Data analytics validates and enriches forecasts, while identifying behavioural trends and the impact of the potential decisions.
- Decision support and DER valorisation integrates these insights into concrete participation strategies and feeds results back into analytics. The results should be also

the inputs for energy sharing or flexibility services mechanisms being defined in 2.3 and 2.4, respectively.

This integrated structure ensures that even when developed and validated as separate proof of concept modules, the methodologies remain aligned toward a common objective: enabling consumers and ECs to benefit from accurate information, context aware recommendations, and transparent and meaningful benefits. This paves the way for the creation of optimised and customised decision-making strategies to improve participation in energy sharing, P2P energy trading, and flexibility mechanisms.

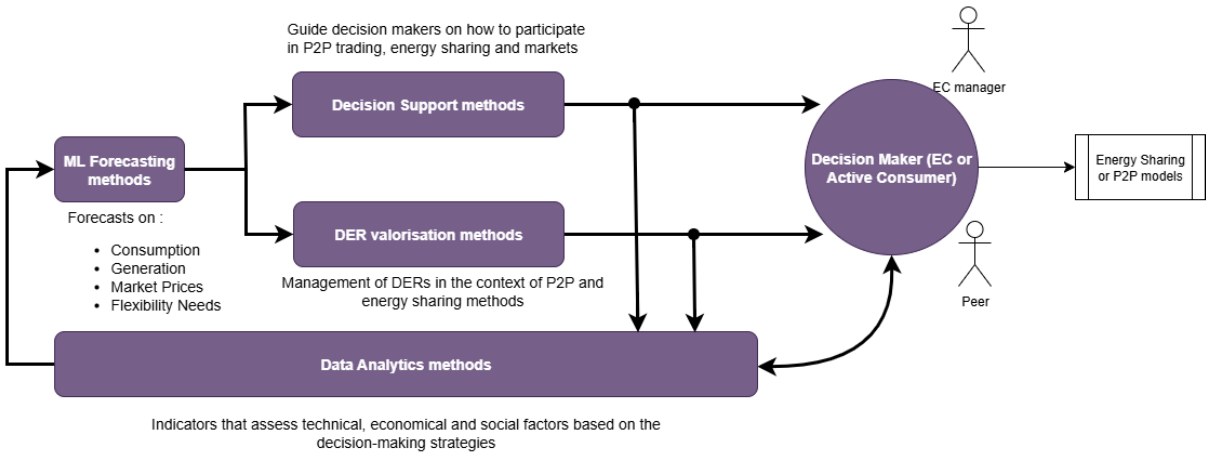

Figure 2-1: Diagram of the framework for the 4 methods under task 2.2

Although applied and demonstrated independently in this deliverable, the methods were designed with complementarity in mind, paving the way for their integration and TRL increase in Task 4.2 within the U2Demo platform. Together, they form the basis for optimised decision-making strategies that enhance participation in data sharing, P2P trading, and flexibility services.

In the following sections each of the four methods are presented in detail as well as their demonstration using data either synthetic or from public available datasets.

3 Machine-Learning based Forecasting methods

3.1 Global objective

The main objective of this section is to provide robust forecasting capabilities that directly support decision-making and analytics for active consumers and energy communities within the U2Demo project. Accurate predictions of energy demand, renewable generation, EV charging loads, and market prices are essential for optimizing participation in P2P energy sharing platforms. By anticipating future scenarios, stakeholders can make informed choices about energy consumption, flexibility provision, and market engagement, ultimately advancing the democratization of energy systems.

Machine learning-based forecasting methods are at the core of this approach, offering the ability to learn complex temporal patterns and adapt to dynamic environments. These methods enable the integration of diverse data sources, such as weather, holidays, and engineered time features, providing a richer context for future decision support. The overarching goal is to move beyond traditional, static models and deliver actionable, context-aware forecasts that empower users to respond proactively to changing conditions.

3.2 Methodology

Section 3.2 focuses on describing the methodology in detail. It introduces the fundamental concepts underlying time series forecasting, outlining the main principles and techniques used to model temporal data. Additionally, this section explains the different evaluation metrics employed to assess the accuracy and performance of the forecasting models, providing a foundation for the analysis presented in the subsequent sections.

3.2.1 Time series Forecasting

Time series forecasting is the process of predicting future values based on previously observed data points that are ordered in time. It is commonly used to estimate trends, seasonal patterns, and other temporal dynamics in fields such as finance, energy, and economics.

3.2.1.1. Lookback window and forecast horizon

Effective forecasting requires defining two crucial parameters: the lookback window and the forecast horizon (**Figure *3-1***).

The lookback window defines the number of past time steps used as input. It determines how much historical context is fed into the model to infer future values.

Let L denote the lookback length. The input to the model at time t is the sequence:

$$X_t = \{y_{t-L}, y_{t-L+1}, \dots, y_{t-1}\} \quad (3\text{-}1)$$

Choosing L involves a trade-off: longer windows may capture more informative patterns but increase computational cost and risk of overfitting. The choice of L is also guided by domain knowledge such as known seasonality and by model evaluation during experiments.

The forecast horizon specifies how far ahead predictions are made. Let H denote the horizon. The model outputs:

$$Y_t = \{y_t, y_{t+1}, \dots, y_{t+H-1}\} \tag{3-2}$$

The forecast horizon is chosen to match the application's requirements.

3.2.1.2. Autoregressive and sequence-to-sequence approaches

When forecasting with a horizon greater than one time step (i.e. multi-step forecasting), two main approaches are commonly used: the autoregressive (AR) strategy, which predicts one step at a time and recursively feeds its own previous predictions as inputs for subsequent steps; and the sequence-to-sequence (seq2seq) approach, which directly forecasts the entire horizon H in a single forward pass, treating the task as a sequence generation problem.

The autoregressive method can accumulate error over time due to its dependence on prior predictions, whereas seq2seq models aim to mitigate this by modelling dependencies across the entire forecast window. However, seq2seq models may be more difficult to train and may underperform when data is scarce.

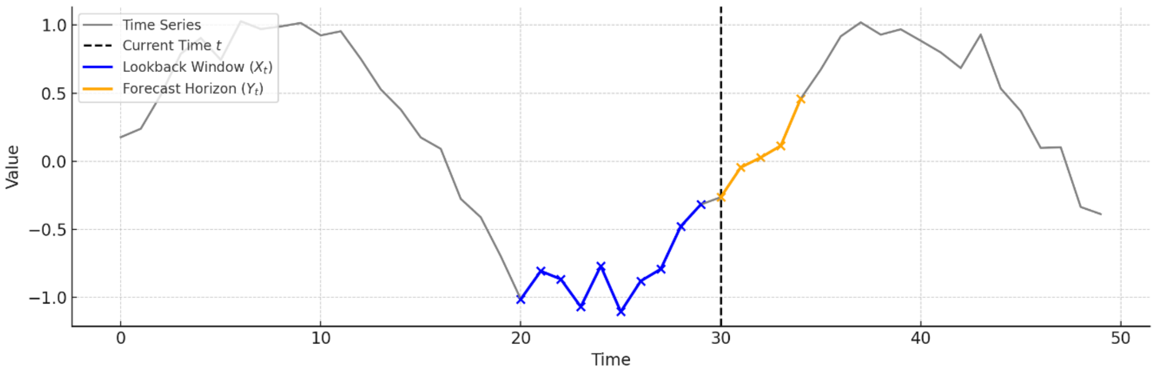

Figure 3-1: Example of lookback window (L=10) and forecast horizon (H=5)

3.2.1.3. Multivariate forecasting

While univariate forecasting can be performed solely using historical data of the target variable (e.g. past electricity consumption or power generation), incorporating additional external information often leads to improved predictive performance.

Multivariate forecasting leverages additional time-varying or static features, known as covariates, are included to capture relevant external influences. For instance, in energy demand forecasting, covariates such as temperature, humidity, wind speed, day of the week, or public holidays are often strong predictors of consumption behaviour.

Covariates are further classified based on their availability at prediction time: known covariates (e.g. time-of-day or holiday schedules) are fully determined in advance, whereas unknown covariates (e.g. unplanned events) are only available for past time steps. The inclusion and proper treatment of these variables are essential for achieving high forecasting accuracy.

3.2.2 Evaluation Metrics

The **Mean Absolute Error** (MAE) measures the average magnitude of the errors in a set of forecasts, without considering their direction. It is defined as:

$$\mathrm{MAE} = \frac{1}{n}\sum_{t=1}^{n}|y_t - \hat{y}_t| \tag{3-3}$$

where y_t is the actual value at time t and $\hat{y}_t$ is the forecasted value, and n is the total number of time steps. MAE is intuitive and easy to interpret, but it does not penalize larger errors more heavily than smaller ones.

The **Weighted Absolute Percentage Error** (WAPE) is a commonly used metric in time series forecasting, particularly in domains like demand forecasting and inventory management. It measures the total absolute forecast error relative to the total actual values and is defined as:

$$\mathrm{WAPE} = \frac{\sum_{t=1}^{n}|y_t - \hat{y}_t|}{\sum_{t=1}^{n} y_t} \times 100\% \tag{3-4}$$

This formulation expresses the cumulative forecast error as a percentage of the total observed values, making it easy to interpret and communicate in practical contexts. WAPE is especially useful in comparing model performance across different time series, as it is scale-independent, assuming the total actual value is meaningful and non-zero.

However, it may become unstable or misleading when the total actual value is very small or zero, and it implicitly assumes that all forecasted points have equal importance.

The **Root Mean Squared Error** (RMSE) calculates the average of the squared differences between forecasted and actual values:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2} \tag{3-5}$$

RMSE places greater emphasis on large errors due to the squaring term, making it more sensitive to outliers. It is particularly useful when large forecasting errors are especially undesirable.

The **Weighted Quantile Loss** (WQL) evaluates the accuracy of probabilistic forecasts at specified quantile levels. It extends the quantile loss by weighting contributions according to the magnitude of the actual values, so that time series or time steps with larger scales have greater influence. WQL is defined as:

$$\mathrm{WQL} = \frac{\sum_{t=1}^{n}\sum_{q\in Q}\rho_q\left(y_t, \hat{y}_t^{(q)}\right)}{\sum_{t=1}^{n}|y_t|} \tag{3-6}$$

where y_t is the actual value at time t, $\hat{y}_t^{(q)}$ is the forecasted value for quantile q, Q is the set of quantiles being evaluated, and ρ_q is the quantile loss function:

$$\rho_q(y, \hat{y}) = q \cdot (y - \hat{y}) \cdot 1_{y \geq \hat{y}} + (1 - q) \cdot (\hat{y} - y) \cdot 1_{y < \hat{y}} \tag{3-7}$$

This formulation makes WQL scale-dependent, as larger actual values increase the denominator and thus influence the metric more. When $Q = \{0.5\}$ (the median), WQL reduces to the WAPE. Unlike point forecast metrics such as MAE or RMSE, WQL measures the calibration and accuracy of full predictive distributions, making it especially useful in probabilistic forecasting scenarios.

In this study we use WQL including every decile ($Q = \{0.1, 0.2, \ldots, 0.9\}$) as the loss function to train machine learning and deep learning models.

3.3 Explored models

We explore a range of forecasting models, spanning simple baseline approaches, classical statistical models, machine learning methods, deep learning architectures, and time series foundation models. These models vary in complexity, interpretability, and scalability. In addition, ensemble methods combine the strengths of individual models to enhance robustness and accuracy.

3.3.1 Baseline Models

Baseline models provide simple benchmarks against which more sophisticated models can be evaluated. They are quick to implement, interpretable, and capture basic persistence or seasonality patterns.

3.3.1.1. Naïve Model

The Naïve Model sets the forecast equal to the last observed value. Quantiles are derived assuming zero-mean residuals with variance estimated empirically. This persistence-based approach provides a simple, reliable benchmark for evaluating more advanced models. [1]

3.3.1.2. Seasonal Naïve Model

The Seasonal Naïve Model predicts future values by using the observed value from the same hour of the previous day. It effectively captures daily seasonal patterns and provides a baseline for comparing other forecasting models. [1]

3.3.1.3. Average Model

The AverageModel predicts the historical mean for future values, offering a simple benchmark for stable series without trends or seasonality. [1]

3.3.1.4. Seasonal Average Model

The Seasonal Average Model computes forecasts as the historical average for the same season, accounting for periodic patterns and providing a straightforward seasonal benchmark. [1]

3.3.2 Statistical Models

Statistical models capture simple, interpretable patterns in the data, such as trends and seasonality, using parametric assumptions. They are efficient, suitable for datasets with clear structures, and often provide analytic confidence intervals, enhancing reliability in practical applications.

3.3.2.1. ETS Model

Exponential Smoothing State Space (ETS) Models extend classical exponential smoothing by explicitly modeling error, trend, and seasonal components. They support additive or multiplicative structures, making them adaptable to a wide range of series with varying growth rates and seasonal patterns. ETS models are estimated with maximum likelihood and can generate both point forecasts and prediction intervals. [2]

3.3.2.2. Theta Model

The Theta method decomposes a time series into components with different curvature levels, combining linear regression for long-term trends with exponential smoothing for short-term dynamics. Its balance of simplicity and accuracy made it highly successful in the M3 forecasting[2] competition. [3]

3.3.3 Machine Learning Models

Classical machine learning models forecast future values using structured features derived from past observations, often framing the problem as regression on engineered covariates. These models are efficient and interpretable.

3.3.3.1. Light Gradient Boosting Machine (LightGBM)

Light Gradient Boosting Machine (LightGBM) is a gradient-boosted decision tree model that forecasts the entire horizon at once (seq2seq). It builds decision trees using a histogram-based algorithm and a leaf-wise tree growth strategy, which significantly speeds up training and reduces memory usage compared to traditional level-wise approaches. [4]

3.3.4 Deep Learning Models

Deep learning models use neural networks to capture complex temporal patterns and non-linear relationships in time series data. They typically require larger datasets to perform effectively but often achieve state-of-the-art (SOTA) results. These models scale well to multiple series, can incorporate covariates, and provide probabilistic forecasts. The most effective architectures for time series forecasting are based on recurrent neural networks (RNNs) or transformers.

[2] Makridakis Competitions (also known as the M Competitions)

3.3.4.1. DeepAR

DeepAR is a probabilistic forecasting model introduced in [5]. It uses an autoregressive RNN architecture, typically based on Long-Short Term Memory (LSTM) architecture, to learn temporal dependencies in univariate and multivariate time series. By modelling the conditional distribution of future values given the past, DeepAR produces probabilistic forecasts rather than point estimates. This makes it particularly effective in settings with high uncertainty and large numbers of related time series. However, due to its sequential nature, DeepAR can be computationally expensive for very long sequences, and struggles to capture long-range dependencies efficiently compared to newer transformer-based models.

3.3.4.2. Temporal Fusion Transformer (TFT)

The Temporal Fusion Transformer (TFT), proposed in [6], is a hybrid deep learning model designed for interpretable multi-horizon forecasting. TFT combines LSTM for local sequence processing with attention mechanisms for capturing long-term dependencies and variable selection. A key contribution of TFT is its interpretability, achieved through variable selection networks and attention weights that highlight influential features across time. This makes TFT particularly useful in applications where understanding drivers of forecasts is as important as accuracy, such as finance and healthcare. Nonetheless, the model's complexity and reliance on recurrent layers can pose scalability challenges for very large datasets.

3.3.4.3. PatchTST

PatchTST, introduced in [7], is a transformer-based model tailored for time series forecasting. Inspired by advances in computer vision, PatchTST represents time series as sequences of patches, enabling the model to capture both local and global temporal patterns more effectively. Unlike recurrent architectures, PatchTST benefits from parallelizable transformer encoders, improving scalability. The patching mechanism not only reduces sequence length but also enhances the model's ability to learn contextual information. PatchTST has demonstrated strong performance across diverse forecasting benchmarks, establishing it as a state-of-the-art method for long-term time series forecasting.

3.3.4.4. Time-Series Dense Encoder

The Time-series Dense Encoder (TiDE), proposed in [8], takes a different approach by employing a feed-forward encoder-decoder architecture rather than recurrent or transformer layers. TiDE is designed for efficiency, offering faster training and inference while maintaining competitive accuracy. The model learns compressed latent representations of time series through dense layers, and its decoder maps these representations to multi-step forecasts. By avoiding sequential dependencies in computation, TiDE scales well to large datasets and long forecasting horizons, making it attractive for industrial applications requiring high throughput.

3.3.5 Foundation Models

Foundation models are large models pre-trained on massive and diverse time series datasets. They enable zero-shot forecasting but can also be fine-tuned for specific applications, offering scalability and generalization capabilities beyond traditional methods.

3.3.5.1. Chronos-Bolt

Chronos-Bolt is a time series foundation model built upon the T5 encoder-decoder transformer architecture and trained on nearly 100 billion diverse time series observations.
It comes in different size variants from 9M to 205M parameters. It can be used off the shelf with zero-shot forecasting or fine-tuned on a specific use case.
In our experiments, we test both Chronos-Bolt base (205M) in zero-shot and Chronos-Bolt small (48M) fine-tuned on our energy time series cases. [9]

3.3.5.2. Weighted Ensemble

In addition to individual forecasting models, weighted ensemble methods provide a strategy to combine the strengths of multiple predictors. By assigning weights to models based on their validation performance or uncertainty estimates, an ensemble can reduce variance and improve robustness compared to relying on a single model. For example, combining probabilistic models such as DeepAR with transformer-based architectures like PatchTST allows the ensemble to balance uncertainty modelling with long-range dependency capture. Weighted ensembles are particularly valuable in heterogeneous forecasting tasks, where no single model consistently outperforms others across all the time series or forecasting horizons.

3.4 Feature engineering

Feature engineering is the process of transforming raw data into meaningful input variables, or "features", that improve the performance of machine learning models. In time series forecasting, feature engineering often involves creating new variables from historical data, such as lagged values, rolling averages, encoded time indicators (like day of the week or hour of the day), and incorporating external factors such as weather or holidays. By carefully designing and selecting features, data scientists help models better capture patterns, trends, and relationships in the data, leading to more accurate and robust predictions. In this section we explore the different features encoded by the different algorithms.

3.4.1 Encoded time features

In time series forecasting, incorporating time-related information as exogenous variables is critical for capturing underlying seasonal patterns, trends, and periodic fluctuations. These time features act as structured prior, enabling models to recognize regularities that are intrinsic to human behavior and environmental cycles.

To effectively utilize temporal structure, time features must represent the cyclic nature of time. For example, the progression from Sunday to Monday or December to January should not be treated as linear steps but as part of a repeating cycle. Linear encoding schemes often fail to capture this periodicity, potentially limiting model performance. Instead, specialized encodings are applied to preserve cyclic properties and enhance model interpretability and predictive power.

In practice, we employ the following encoding strategies:

- **One-hot encoding for day of the week:** This categorical encoding allows the model to treat each weekday independently without implying ordinal relationships. It helps capture behavioral patterns such as increased consumption on weekdays versus weekends.

- **B-spline encoding for day of the year:** B-splines (Figure) offer a smooth and flexible representation of seasonal trends, ideal for capturing gradual changes throughout the year. This is especially relevant for modeling annual cycles influenced by temperature, daylight hours, or calendar events.
- **Sine and cosine encoding for the month:** To represent the cyclical nature of the calendar year, we map the month onto a unit circle using sine and cosine transformations. This allows the model to recognize, for example, that January and December are adjacent in time despite being at opposite ends of a linear scale.
- **Sine and cosine encoding for hour of the day:** Similarly, hourly patterns are encoded using sine and cosine functions to reflect their 24-hour periodicity (Figure).

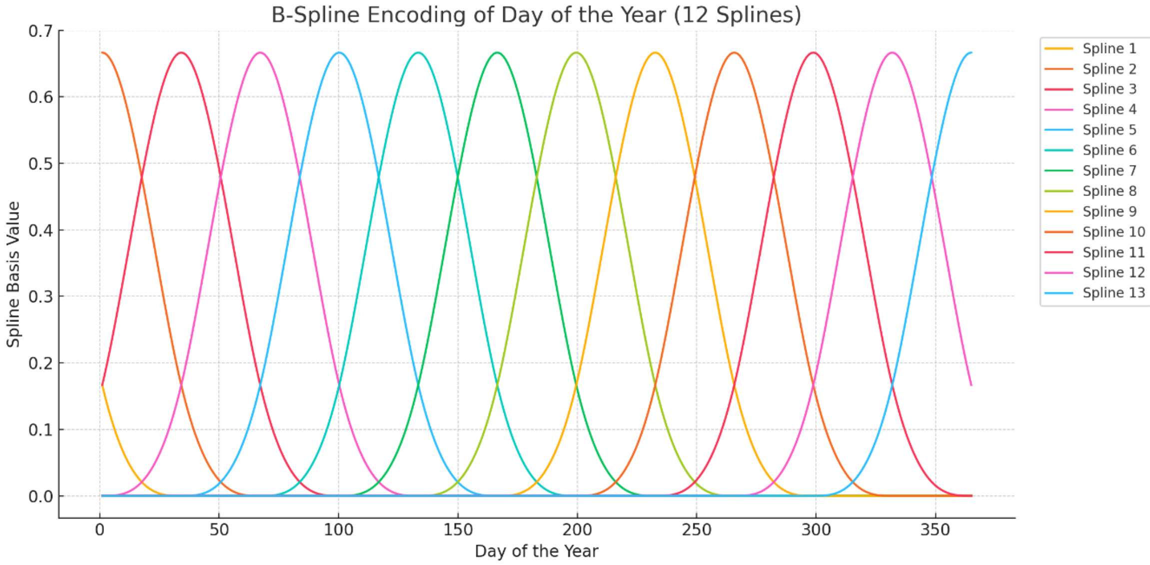

Figure 3-2: Graph representation of B-spline encoding for the day of the year

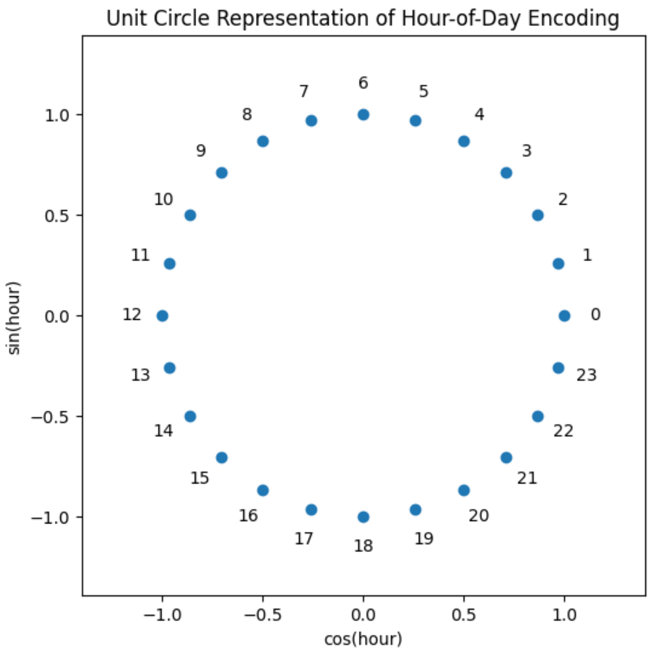

Figure 3-3: Graph representation of sine cosine encoding for the hour of the day

By incorporating these encoded time features, forecasting models can more effectively learn temporal dependencies and align their predictions with known seasonal structures. This is particularly beneficial in energy applications, where demand and generation exhibit strong and predictable temporal rhythms.

3.4.2 Holidays features

In addition to continuous and cyclic temporal features, calendar-based events such as holidays serve as important exogenous variables for time series forecasting. Holidays often induce atypical patterns in energy consumption due to changes in human behavior, industrial activity, and mobility. For example, residential electricity usage may spike during national holidays, resulting in temporary deviations from regular trends.

Unlike cyclic time features, holidays are non-periodic and often irregular, requiring explicit representation rather than transformation. We encode holidays as binary indicator variables, where each holiday (or group of holidays) is assigned, a dedicated variable denoting its occurrence. These indicators can be country- or region-specific, capturing localized effects that global time features may overlook.

By incorporating holiday indicators, models gain the ability to adjust forecasts during known disruptions, improving both accuracy and robustness in operational energy systems.

3.4.3 Weather features

Weather data constitutes a critical class of exogenous variables in time series forecasting, particularly in the context of energy systems. In energy consumption, weather influences human behavior and building operations such as driving heating, ventilation, and air conditioning (HVAC) usage in response to temperature, humidity, and other environmental conditions. In renewable energy production, weather has a direct physical impact: solar irradiance governs photovoltaic output, wind speed affects wind turbine generation, and ambient conditions influence efficiency and intermittency.

To capture these dynamics, historical meteorological variables such as temperature, dew point, wind speed, solar irradiance, and relative humidity are often included as lagged features or aligned with the target timestamp. These variables introduce critical information about the environmental drivers of energy demand and supply, enabling the model to respond more accurately to variability induced by external conditions.

Moreover, weather forecasts can be integrated as known future covariates. This is particularly advantageous in short- to medium-term forecasting, where accurate meteorological predictions from external sources can enhance the foresight of models, especially in renewable energy forecasting. For example, incorporating forecasted solar irradiance or wind speed allows models to anticipate fluctuations in generation capacity before they occur.

By leveraging both historical and forecasted weather data, models gain access to exogenous drivers that are otherwise difficult to infer from the target series alone. This leads to substantial improvements in both accuracy and responsiveness, especially in settings characterized by high environmental volatility or renewable integration.

For this study, we access historical hourly weather data using Open Meteo API [10]. As archived weather forecasts are not easily available for older dates, we also use historical measured weather data for known future weather. This can lead to an overestimation of models' evaluation compared to real-case scenario as the future weather inputs do not contain forecast errors.

3.4.4 Lagged features and rolling statistics

In addition to exogenous variables, a time series' own historical values often provide the strongest predictive signal for future observations. By past values of the target variable shifted by a fixed number of time steps, models can directly capture autocorrelation structures inherent in the data. For example, in hourly energy demand forecasting, the demand observed exactly 24 hours earlier (lag = 24) may serve as a strong predictor of current demand due to daily repetition in consumption patterns.

Lagged features can be constructed at multiple temporal offsets to represent different types of dependencies, such as short-term persistence (e.g., lag 1–3) and longer-term seasonal effects (e.g., lag 24, lag 168 for daily and weekly periodicity in hourly data). Care must be taken to avoid excessive lag depth, as it increases dimensionality and may introduce irrelevant noise.

Complementary to lagged features, rolling statistics summarize the recent behavior of the target variable over a fixed-size window. These statistics include:

- **Rolling mean:** smooths short-term fluctuations and highlights underlying trends.
- **Rolling standard deviation:** measures recent volatility, capturing variability changes that may precede shifts in the series.
- **Rolling minimum/maximum:** identifies extreme values within the recent past, useful for systems sensitive to operational thresholds.

By providing the model with aggregated historical context, rolling statistics help capture dynamics that are not immediately visible from single-point lags. For instance, an elevated rolling mean over the past week may signal a sustained change in baseline demand, while an increasing rolling standard deviation could indicate a period of instability or high variability in renewable generation.

When engineering lagged and rolling features, it is essential to ensure causal integrity: all transformations must use only information available up to the prediction time, avoiding any data leakage from the future.

The integration of lagged target features and rolling statistics enriches the feature space with direct representations of the series' own memory and variability. In combination with exogenous inputs, these engineered features enable models to capture both internal temporal dependencies and external drivers, enhancing short-term reactivity and long-term pattern recognition.

3.5 Forecasting scenarios

Since data from the pilot sites were not available by the time the Task 2.2 was performed, multiple open datasets were selected for each forecasting scenario as alternative experiment ground. They were selected based on their availability and closeness to the expected pilot sites configurations.

3.6 Demand forecasting

In this section, various approaches to demand forecasting are examined. The implemented algorithms are specifically designed to predict the aggregated energy consumption of a household. Forecasting at the level of individual appliances, such as heating systems or washing machines, is beyond the scope of this study. Instead, the prediction focuses exclusively on the total household consumption profile.

3.6.1 Dataset

The dataset used for household electricity load forecasting is the UCI Individual Household Electric Power Consumption dataset [11]. It records electricity consumption of a single household in Sceaux, France, between December 2006 and November 2010, sampled at a one-minute resolution.

The dataset contains over 2 million observations and several variables describing the household's electricity use: Global_active_power (kW), Global_reactive_power (kW), Voltage (V), Global_intensity (A), and three sub-metering channels (kitchen, laundry, heating/cooling). In this study, we focus on Global_active_power as the main target variable. The series is converted from kW to kWh per minute to better align with energy usage forecasting. [11]

3.6.2 Data treatment

The raw dataset is cleaned by removing invalid or missing values and duplicate timestamps. Only Global_active_power is retained as the target variable, and the data is resampled to an hourly resolution since minute-level data is not expected in T4.2. Temporal features (including holidays), weather information, and lagged versions of the target at multiple seasonal and rolling windows are added as exogenous covariates.

Data is split using sliding windows, we set aside the 5 last weeks of data from each site, 4 weeks as validation data and the last week as test data to avoid any data leak. The rest of the data is used as training data for the models.

3.6.3 Method comparison

Household electricity consumption exhibits high variability, reflecting appliance use, daily routines, and weather-dependent patterns. To understand the factors that drive accurate forecasts, we focus on the role of covariates.

Eight covariate configurations are evaluated: univariate (no covariates), temporal only, weather only, lag only, temporal + weather, temporal + lag, weather + lag, and all features. This systematic comparison allows us to isolate the contribution of each type of information. Temporal covariates capture seasonality and calendar effects, weather covariates reflect external conditions affecting energy use, and lag covariates leverage the autocorrelation inherent in consumption patterns.

For each configuration, models are trained using varying lookback windows and forecast horizons (from 1 hour up to 7 days). Forecast accuracy is assessed using rolling test windows, and performance metrics are recorded and compared across runs.

This setup enables a nuanced analysis of feature importance: at short horizons, lag features may dominate due to strong short-term correlations; at daily or weekly horizons, temporal features such as day-of-week and holiday effects gain relevance; weather features can provide complementary value, particularly for heating or cooling demand. Combining all covariates does not necessarily yield the best performance, and part of the study quantifies whether simpler feature sets can achieve comparable accuracy.

By centring the analysis on covariate contribution, we aim to identify which contextual signals are most valuable for household electricity load forecasting and how their relevance varies with forecast horizon and model configuration.

3.6.4 Results

3.6.4.1. Covariate impact on performance

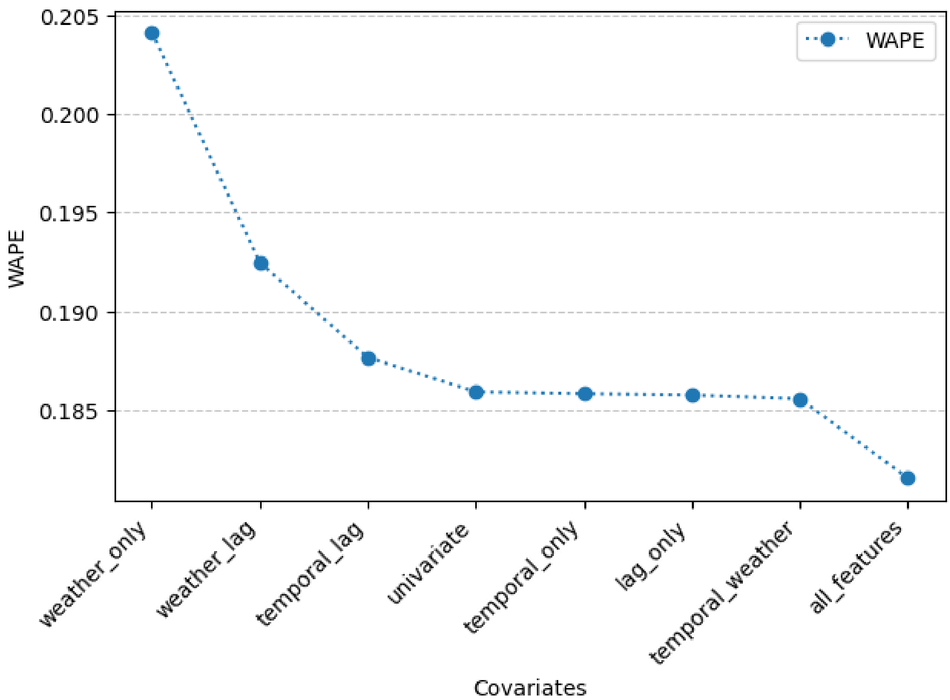

Figure 3-4: illustrates the impact of various covariate combinations on model performance. The x-axis categorizes the covariate configurations, including combinations of lag, temporal, and weather features, as well as univariate baselines.

The results indicate that the model achieves its best performance when all features are combined, yielding the lowest WAPE (**Figure *3-4***). Among individual feature groups, lag features contribute most significantly to performance improvement, while weather-related features, particularly when used in isolation, result in the lowest performance, even worse than the univariate baseline. This suggests that, for the consumption forecasting use case, weather data has limited predictive power when used alone. However, it may still provide complementary value when integrated with other covariates.

It is important to note that the figure reflects the maximum model performance across various lookback and forecast horizon configurations, rather than the mean. This choice was made to avoid the smoothing effect observed in averaged results, which tended to obscure the relative impact of individual covariates.

3.6.4.2. Horizon impact on performance

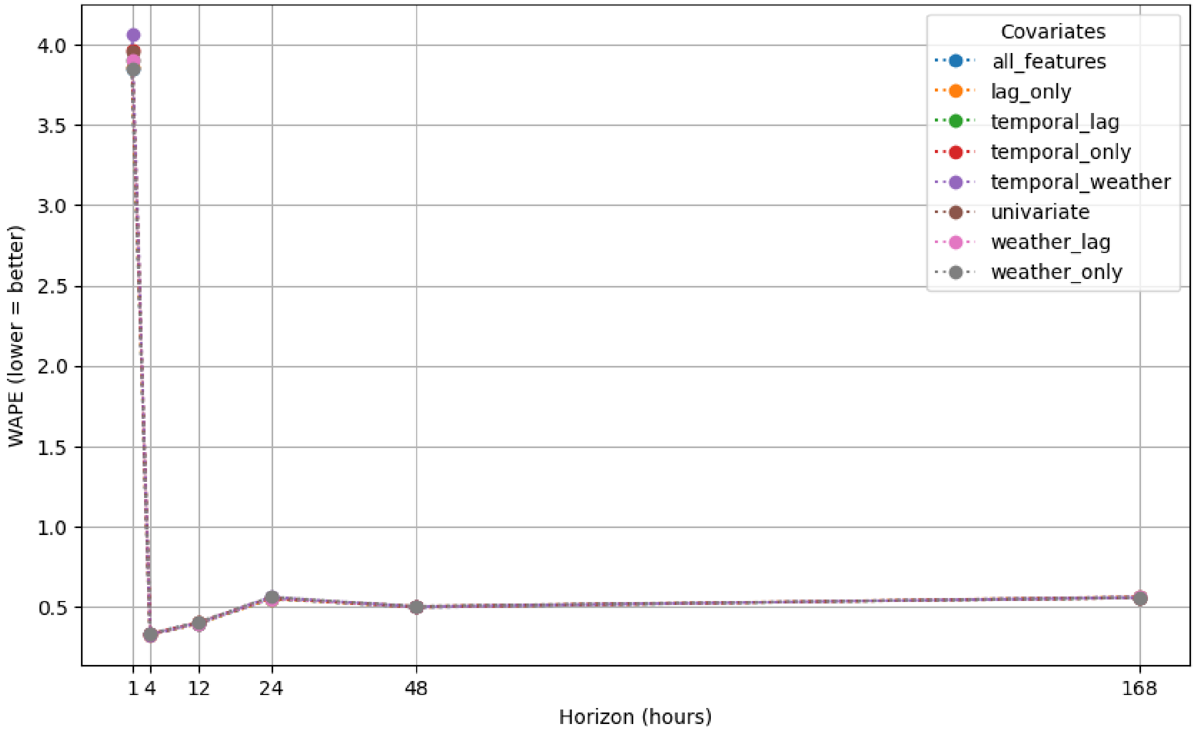

Figure 3-5: illustrates the impact of forecast horizon length on mean model performance, across various covariate configurations. The x-axis represents the forecast horizon in steps (hours).

The results reveal a counterintuitive pattern at the 1-hour horizon, which yields the highest error, suggesting that very short-term forecasts may be more volatile or sensitive to noise.

Performance improves sharply at the 4-hour horizon and remains consistently low through longer horizons up to 168 hours. This stability indicates that once past the short-term range, increasing the forecast horizon does not significantly degrade accuracy.

Notably, the choice of covariates has limited influence on performance across horizons.

3.6.4.3. Training horizon impact on generalization

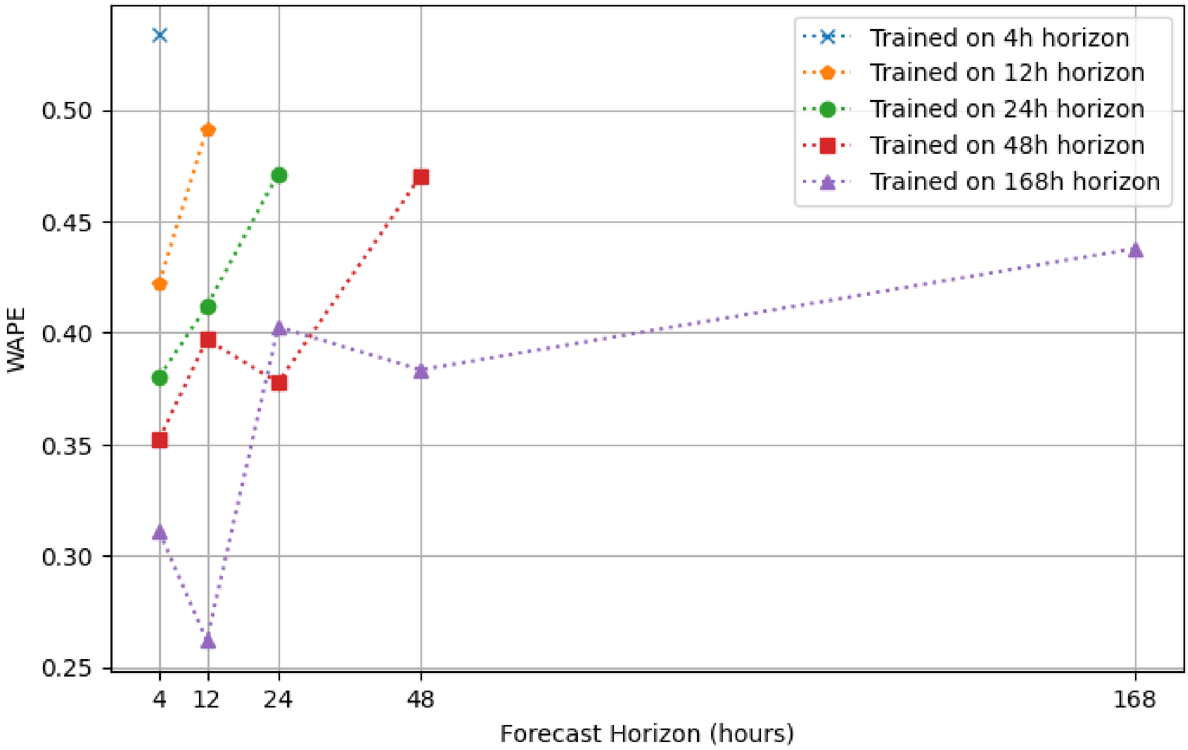

Figure 3-6: explores how the training horizon influences model performance across varying inference horizons. The x-axis represents the forecast horizon in hours, while each line corresponds to a model trained on a specific horizon (4h, 12h, 24h, 48h, or 168h).

The results from **Figure *3-6*** show that models tend to perform better on shorter forecast horizons than the ones they were specifically trained for. For instance, the model trained on a 24-hour horizon exhibits its highest error when forecasting at 24 hours, while achieving lower errors on shorter horizons.

In many cases, these models even outperform those trained directly on the target horizon. Notably, the model trained on the longest horizon (168 hours) performs better in almost every case.

This suggests that training models on longer horizons may be beneficial, as they tend to generalize more effectively across different forecasting intervals.

Table 3-1: Performance of evaluated models trained with a horizon of 48h and all exogenous features on test data.

Model	WAPE@24h (%)	MAE@24h (kW)	RMSE@24h (kW)
WeightedEnsemble	**21.59**	**0.24**	**0.34**
Chronos Bolt Small (Zero Shot)	22.07	0.25	0.34
TemporalFusionTransformer	23.00	0.26	0.38
SeasonalNaive	24.33	0.28	0.39
GBM	30.08	0.34	0.47
Recursive GBM	30.70	0.35	0.45
Theta	46.11	0.52	0.63
ETS	69.40	0.78	0.95
Naive	72.64	0.82	1.00

The best-performing model from Table 3-1 was the Weighted Ensemble, achieving 21.59% WAPE@24h, outperforming all individual models. Interestingly, the Chronos Bolt Small foundation model, evaluated in zero-shot, reached a comparable performance of 22.07% WAPE@24h, showcasing strong generalization capabilities. Among traditional models, the Temporal Fusion Transformer performed best with 23.00% WAPE@24h.

Surprisingly, the Seasonal Naive baseline achieved 24.33% WAPE@24h, outperforming several more complex models such as GBM and Recursive GBM. This suggests that the underlying data may have strong seasonal patterns that are effectively captured even by simple heuristics, highlighting the importance of benchmarking against robust baselines.

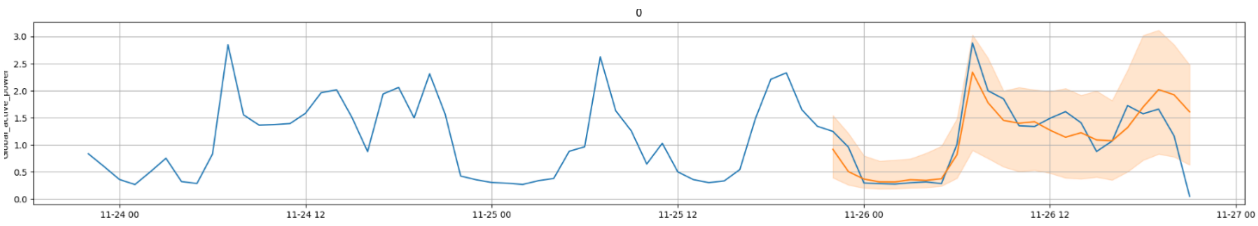

Figure 3-7: Example inference with a horizon of 24h on test data with the best Weighted Ensemble model trained with all covariates features. In blue is the observed data. The orange line is the mean forecasted value, and the orange area represents the 10th-90th percentile interval.

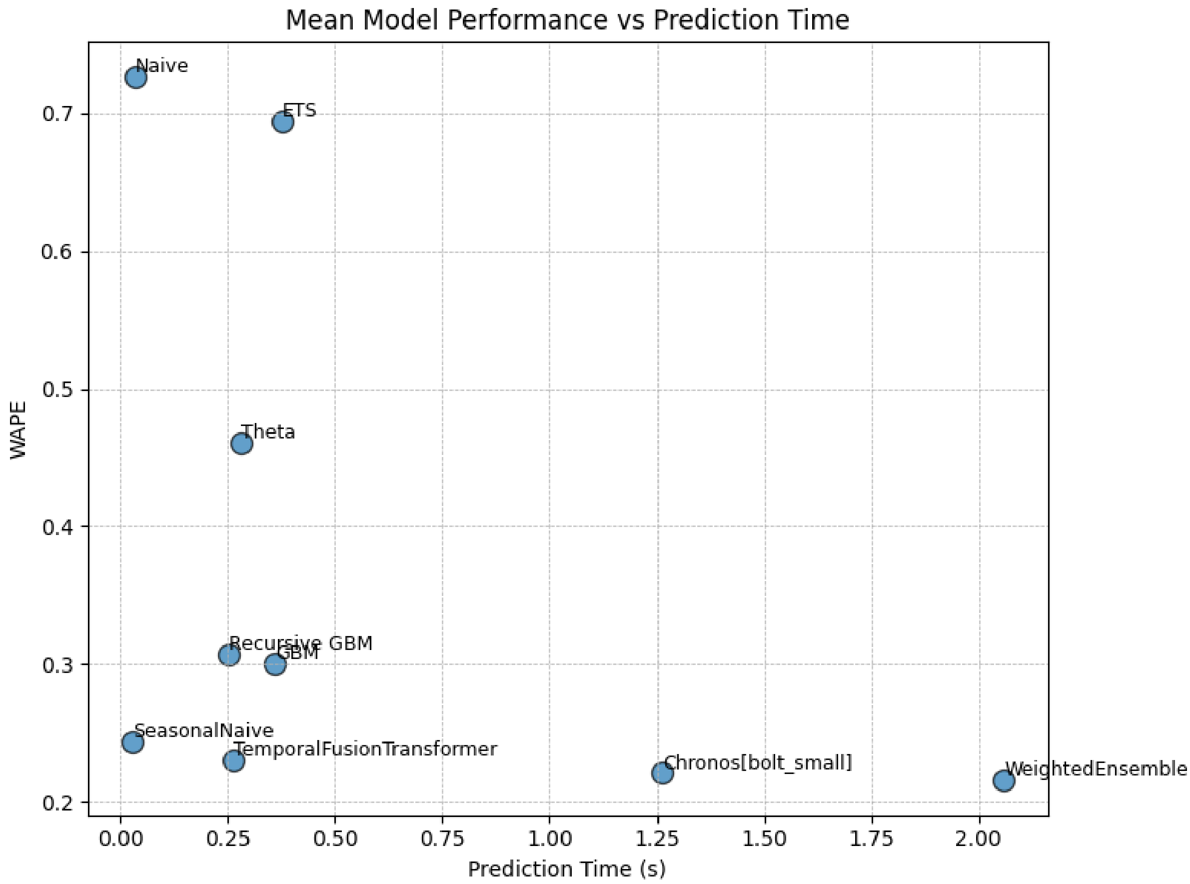

Figure 3-8: Model Accuracy vs Prediction Time in Consumption Forecasting

Lastly, we wanted to illustrate the trade-off between model performance and prediction time across various forecasting approaches (Figure 3-8). Simpler models like Naive, ETS, Theta, and SeasonalNaive predict quickly, but their performance varies with higher WAPE values. Mid-tier models such as Recursive GBM offer moderate speed and improved accuracy. More advanced models, including TemporalFusionTransformer, Chronos[bolt_small], and WeightedEnsemble, generally take longer but achieve lower WAPE, with GBM standing out as an efficient option that delivers the best accuracy at a relatively low prediction time. As expected, models using advanced machine learning or deep learning techniques require more time in most cases. We only highlight prediction time in this use case, as the graph would look similar across other use cases, making further repetition unnecessary. We can also note that prediction time remains relatively low overall, may vary depending on the length of the forecast horizon, and was not a requirement for the project, so the focus remains on the accuracy of the predictions.

3.7 Electrical Vehicle Charging Station load forecasting

In this section, different approaches to load forecasting for electric vehicle (EV) charging stations are examined. The study investigates both the forecasting of individual charging points and the aggregated load of the entire charging infrastructure. The results highlight the importance of data aggregation to obtain more accurate and reliable predictions.

3.8 Dataset

The dataset used for Electrical Vehicle (EV) Charging Station load forecasting is an aggregation of the Public EV Charge Point Usage Dundee City Council datasets [12] spanning from July 2021 to December 2024. It records around 330,000 charging events from 100 public charging points in the city of Dundee, Scotland.

3.8.1 Data treatment

The data is cleaned by removing outlier charging events (e.g. negative consumption during charging or consumption too large). We also remove bus charging station are they represent a particular case that is out-of-scope for our study.

The charging events are then grouped by charging points and aggregated to a given time step (daily or hourly) to transform the raw data into time series data.

Encoded time features, lagged target values, weather data and local holidays are added as exogenous covariates.

Data is split using sliding windows, we set aside the 5 last weeks of data from each site, 4 weeks of as validation data and the last week as test data to avoid any data leak. The rest of the data is used as training data for the models.

3.8.2 Method comparison

EV charging station load time series are highly intermittent. At the level of a single charging point, the load is typically zero for long periods, punctuated by short, high-magnitude peaks corresponding to charging sessions (Figure 3-9). This results in a flat baseline interspersed with sudden, irregular spikes, making accurate forecasting particularly challenging.

To address this issue, we explore aggregation strategies that reduce sparsity and variability in the signal. Two types of aggregation are considered:

- Temporal aggregation – where the time resolution of the data is reduced, e.g., converting hourly measurements into daily averages or totals. This smooths short-term fluctuations and may help the forecasting model capture broader usage patterns, such as daily or weekly trends.
- Spatial aggregation – where the load from multiple charging points is summed to form a single aggregated time series. This approach reduces intermittency by averaging out the stochasticity of individual usage events, leveraging the diversity of charging behaviors across different points.

In this work, we evaluate and compare the performance of these aggregation strategies for EV charging station load forecasting. We examine whether temporal aggregation, spatial aggregation, or a combination of both leads to improved predictive accuracy, and we quantify the trade-offs between resolution, stability, and model performance.

Beyond aggregation strategies, we investigate the impact of different modeling approaches on EV charging station load forecasting performance. We compare a range of classical ML, DL architectures.

We also assess the contribution of exogenous variables in improving forecast accuracy. This analysis provides insight into which contextual signals are most valuable for reducing prediction error and whether their usefulness depends on the chosen modeling approach.

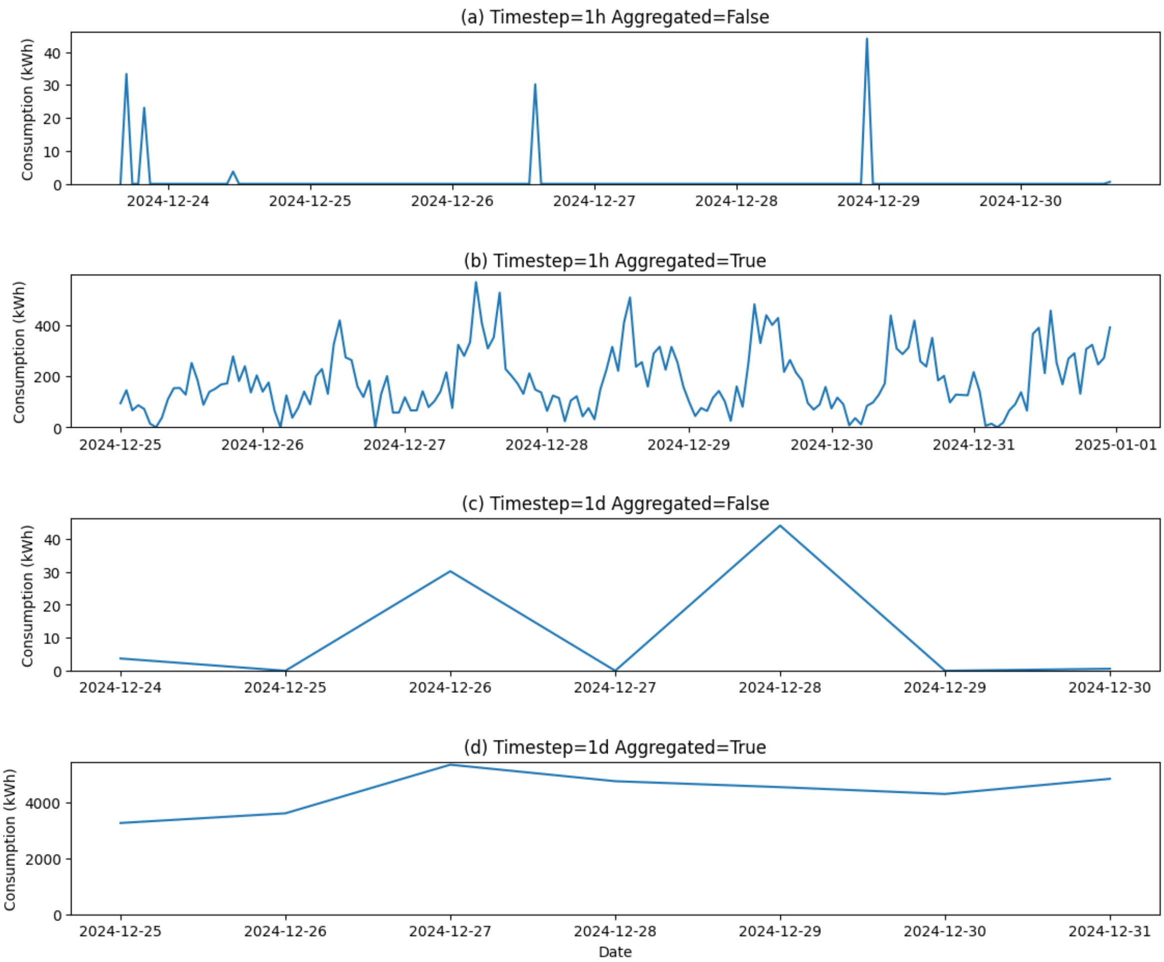

Figure 3-9: Comparison of spatial and temporal aggregation on one week of data. With (a) a timestep of 1 hour without charging stations aggregation, (b) a timestep of 1 hour and aggregated stations, (c) a timestep of 1 day without aggregation and (d) a timestep of 1 day with aggregation.

3.8.3 Results

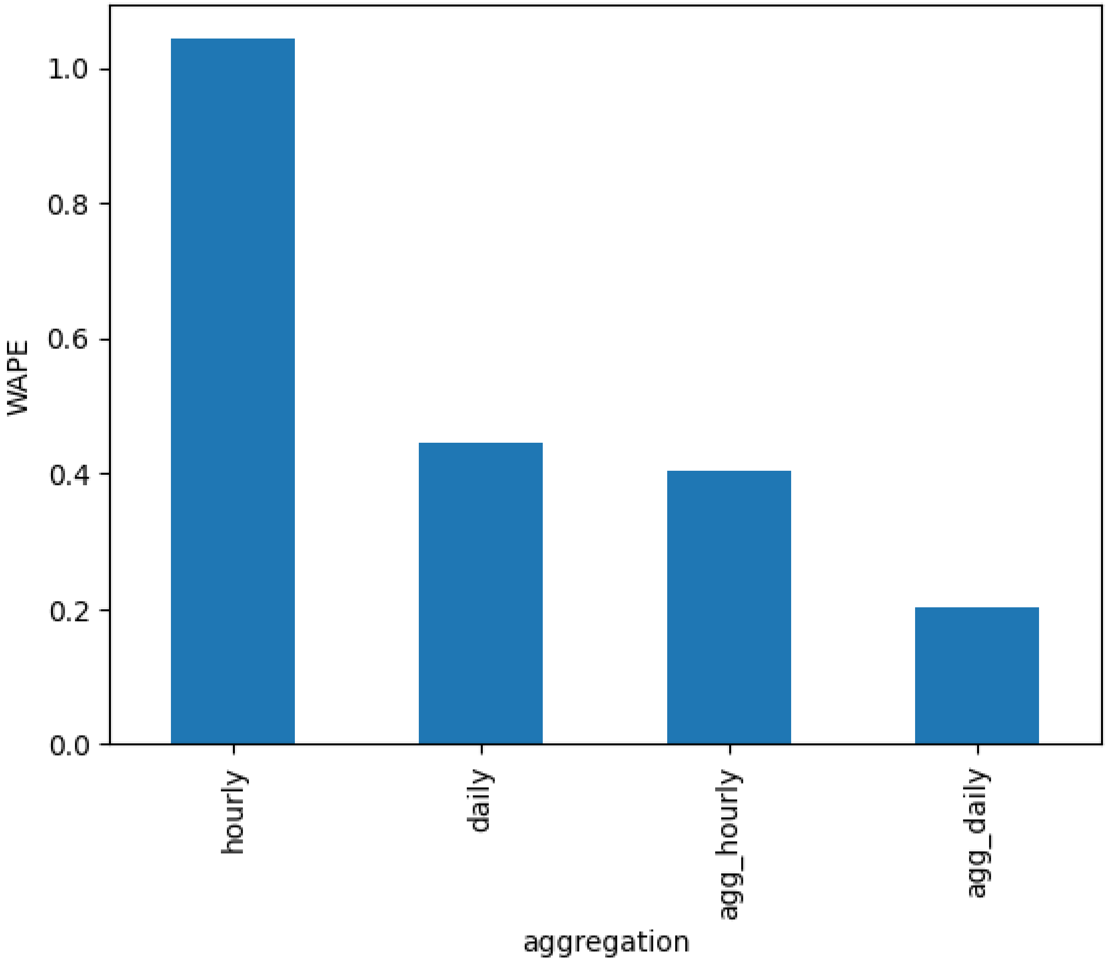

Figure 3-10: Impact of data aggregation on averaged forecasting test error. Hourly and daily stand for the temporal resolution and the "agg_" prefix corresponds to aggregated charging stations.

We observe that by aggregating the data (**Figure *3-10***), we make it more predictable, from more than 100% WAPE to less than 20% when aggregated all the stations to a daily timestep. Although this was the best case, we conduct the following studies with the aggregated stations and hourly timestep as we want to keep a high temporal resolution.

We measure the impact of the additional exogenous covariates on the best performing models (**Figure 3-11**). While PatchTST's best performance is reached with only holidays and past weather data, the DeepAR, GBM and the overall best score are reached with all the exogenous variables included (holidays, encoded time features, lagged features and both past and future weather).

We also study the impact of training and inference forecast horizon on the models' performance (**Figure 3-12**). At first, the error decreases when increasing the forecast horizon then slightly increases when forecasting more than 48h ahead. We also observe that when trained on longer horizons (i.e. more than 96 hours), the models tend to perform better on shorter horizons than models solely trained on those. For example, the models trained on 168h horizons reach 30.6% WAPE@24h on average while the models trained on 24h horizons only achieve 31.7% WAPE@24 on average.

The best single model evaluated in this study was the DeepAR model that reached 25.25% WAPE (**Table 3-2**). It was trained on the data with all the additional features proposed (i.e. holidays, encoded time, lagged target, past and future weather) on a training horizon of 168h and evaluated on an inference horizon of 48h.

The weighted ensemble model slightly improved on this score reaching 25.20% with both training and inference horizons of 48h.

We can also note that the foundation model Chronos Bolt Base evaluated in zero-shot achieves similar performance with 25.32%.

Table 3-2: Performance of evaluated models trained with a horizon of 168h (or 48h when marked with *) and all exogenous features on test data.

Model	WAPE@48h (%)	MAE@48h (kWh)	RMSE@48h (kWh)
Weighted Ensemble*	25.20	62.72	84.00
DeepAR	25.25	62.83	84.07
Chronos Bolt base (Zero Shot)	25.32	63.03	84.16
GBM	25.60	63.60	83.77
TFT	25.65	63.88	84.60
TiDE	26.11	64.98	85.94
PatchTST*	26.52	66.00	86.27
Chronos Bolt small (finetuned)	26.53	66.20	87.81
Seasonal Naive	40.81	102.00	130.79

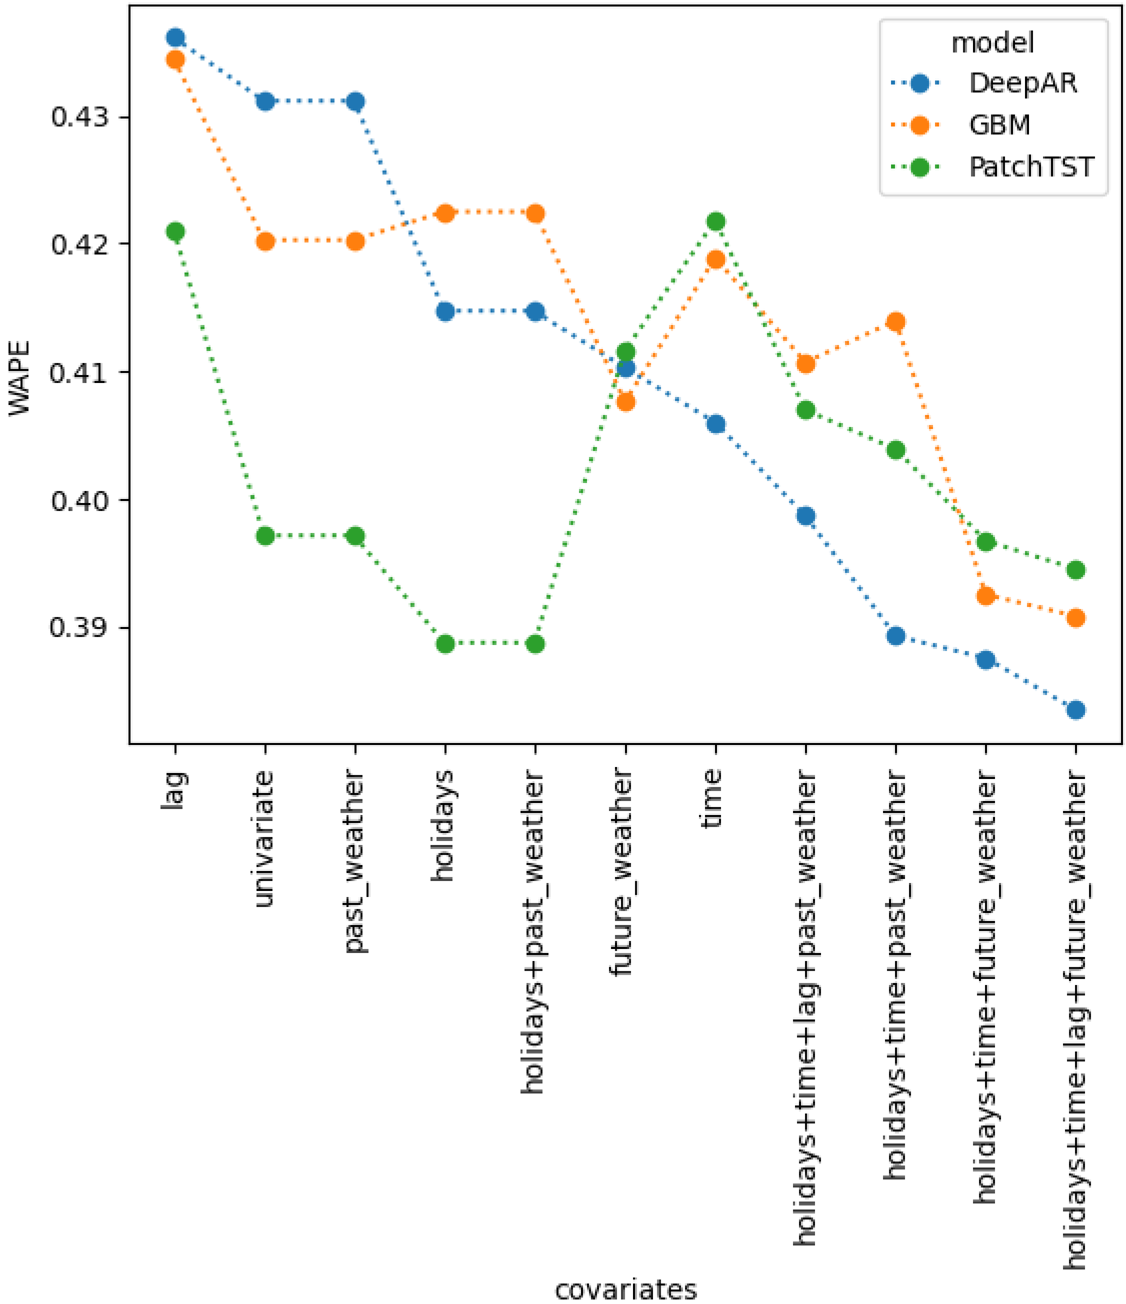

Figure 3-11: Impact of exogenous covariates on forecasting test error.

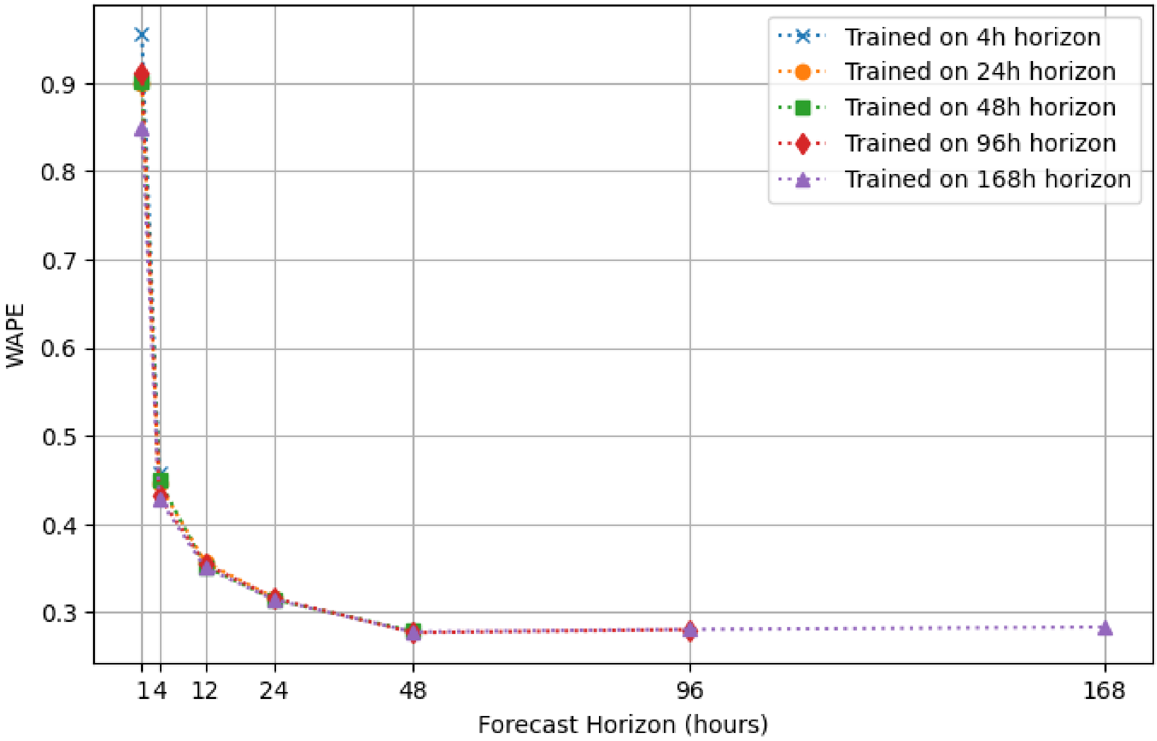

Figure 3-12: Impact of training and inference horizons on the averaged model performance on aggregated EV charging station load forecasting.

3.9 Solar production forecasting

This section benchmarks short-term photovoltaic (PV) generation forecasting methods that support decision making for active consumers and energy. The focus is on hourly forecasts at residential scale, using machine learning approaches that combine past PV production with exogenous drivers, primarily driven by weather conditions and temporal patterns such as hour of day and seasonality

3.9.1 Dataset

The dataset used for PV energy generation forecasting is the Photovoltaic Solar Panel Energy Generation data provided by UK Power Networks [13]. It regroups, 10-minute interval of averaged generated power measurements from 6 domestic sites with solar panels nearby London. The measurements span from June to November 2014 which makes around 131,000 individual measurements.

3.9.2 Data treatment

Only the averaged generated power kW is kept as the target and we convert it to the generated energy in kWh. The data is aggregated to an hourly timestep to smoothen the data and avoid having too much noise. We round up negative energy values to zeros and interpolate small missing value windows.

Encoded time features, lagged target values and weather data are added as exogenous covariates.

The weather data is pulled using Open Meteo API [10] for the city of London, in the case of known future weather we use historical observed data instead of forecasted weather because of availability but we would need to use forecast in real case

Data is split using sliding windows, we set aside the 5 last weeks of data from each site, 4 weeks as validation data and the last week as test data to avoid any data leak. The rest of the data is used as training data for the models.

3.9.3 Method comparison

We first compare different models with a lookback window of 168 hours (corresponding to one week) and a horizon of 24 hours.

We then measure how the error evolves when changing the horizon and also conduct ablation studies on the use of exogenous covariates.

For the use case of PV energy production, we avoid using WAPE metric as it becomes less reliable when dealing with many zero values which is the case of PV production overnight. Instead, we mainly use MAE which is more meaningful and readable.

3.9.4 Results

Table 3-3: Performance of evaluated models trained with a horizon of 24h and future_weather features on test data.

Model	MAE (kWh) ↓	RMSE (kWh)
Weighted Ensemble	0.0792	0.1984
GBM	0.0808	0.2053
DeepAR	0.0812	0.2096
Chronos Bolt base (Zero Shot)	0.0945	0.2247
TiDE	0.0975	0.2373
Seasonal Naïve	0.1112	0.2941
Chronos Bolt small (finetuned)	0.1237	0.2397
PatchTST	0.1476	0.3862
TFT	0.2070	0.3916

When evaluating the impact of horizon on forecasting error (**Figure 3-13**), we observe a significant dip in the error for the 4h horizon and the performance stabilizes when forecasting for more than 24h ahead. We also observe that in this case, models trained for shorter horizons slightly outperform the ones trained on longer horizons

When comparing the impact of exogenous covariates (**Figure 3-15**), we observe that the *future_weather* is most impacting covariate. This is expected as the weather and especially the solar irradiance directly impacts the PV energy production. The other evaluated features have little to no positive impact on the models' performance.

We focus on the forecasts with a horizon of at least 24h as shorter terms forecasts are not as useful for decision-making (**Table 3-3**). In this case the best performing models were GBM trained and evaluated with a horizon of 24h with an MAE of 81Wh and the weighted ensemble model which achieved 79Wh MAE (**Figure 3-14**).

The results and especially the impact of weather data have to be validated for the demo sites. In our study the weather data was retrieved using Open Meteo API [10] for the city of London without taking into account the precise location of the different PV panels sites. A more precise weather data gathering might improve the accuracy of the models. However, the weather data contained in the *future_weath*er feature is also historical weather data, as archived weather forecasts are not easily accessible. The impact of actual weather forecasts as a feature has to be further investigated and may harm the models' performance depending on their quality.

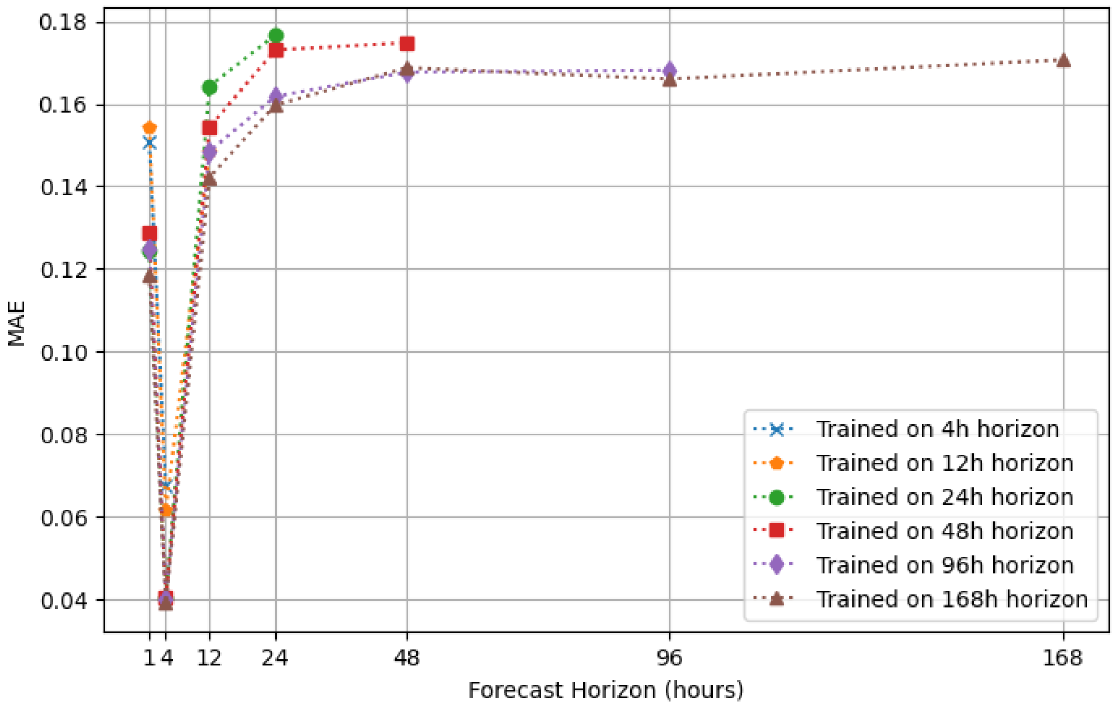

Figure 3-13: Impact of training and inference horizons on the averaged model performance on PV production forecast.

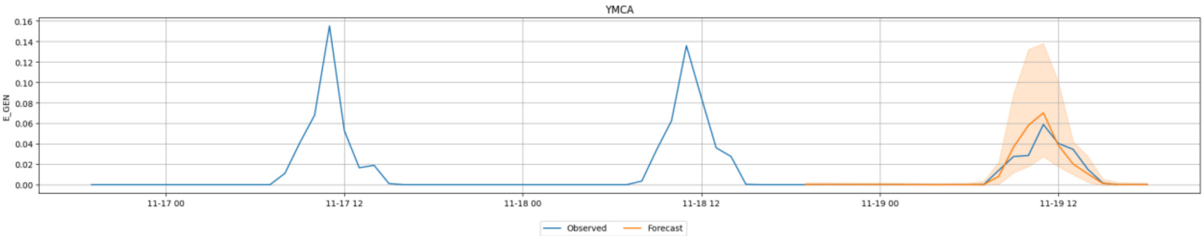

Figure 3-14: Example inference with a horizon of 24h on test data with the best Weighted Ensemble model trained with future_weather feature. In blue is the observed data. The orange line is the mean forecasted value and the orange area represents the 10th-90th percentile interval.

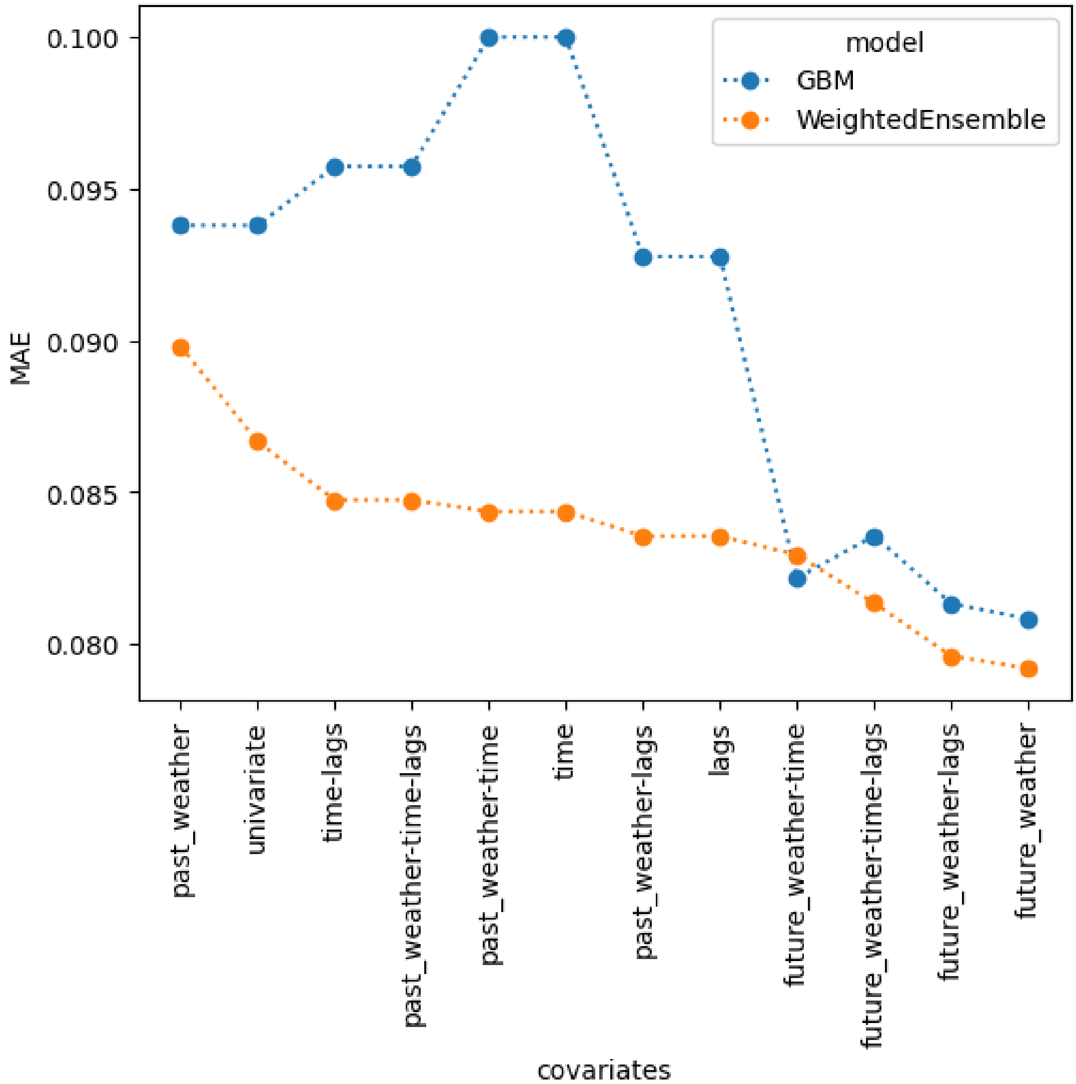

Figure 3-15: Impact of exogenous covariates on forecasting test error of the best performing models.

3.10 Wind production forecasting

This section explores short-term wind generation forecasting methods to support decision making for active consumers and energy communities. Although the current U2Demo pilot sites do not operate wind assets, we include this study to ensure replicability of the forecasting framework in communities that do, and to ensure the methodology remains consistent across all energy vectors considered in U2Demo, including consumption, PV, and wind, for comprehensive strategy and flexibility planning.

3.10.1 Dataset

For wind production forecasting, we use an open dataset [14] of historical hourly information of four European wind farms located in France, Italy and Spain spanning from April 2021 to November 2023.

3.10.2 Data treatment

Some of the time series have long windows of missing values so we split them in multiple separate consistent time series, and we fill the remaining smaller windows using interpolation.

Encoded time features, lagged target values and weather data are added as exogenous covariates.

Data is split using sliding windows, we set aside the 5 last weeks of data from each site, 4 weeks of as validation data and the last week as test data to avoid any data leak. The rest of the data is used as training data for the models.

3.10.3 Method comparison

We follow a similar methodology as for the solar production case.

First, we compare different models with a lookback window of 168 hours (corresponding to one week) and a horizon of 24 hours.

We then measure how the error evolves when changing the horizon and also conduct ablation studies on the use of exogenous covariates.

3.10.4 Results

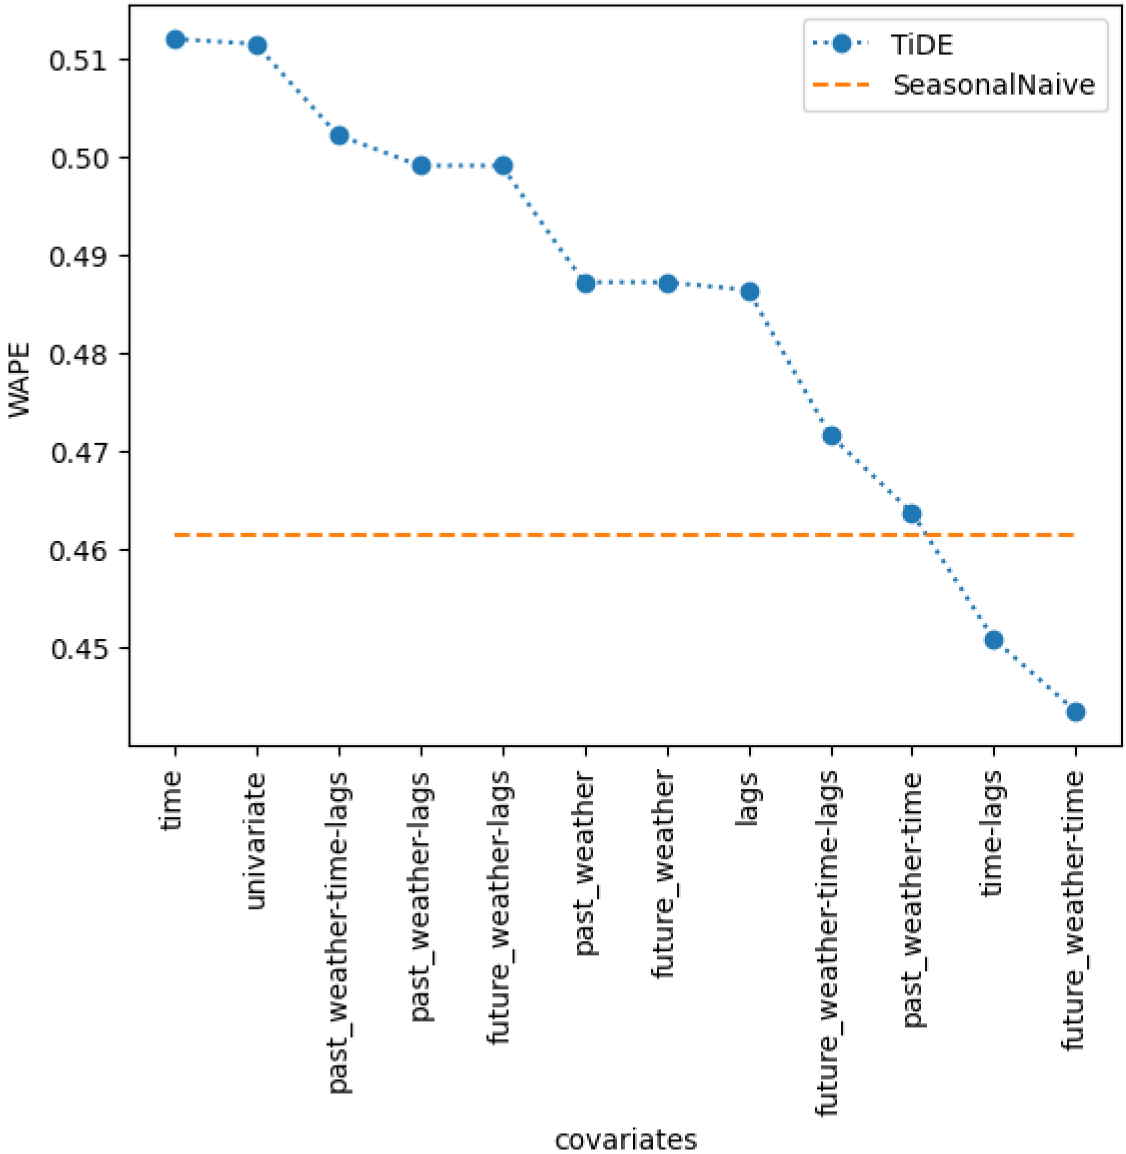

Figure 3-16: Impact of exogenous covariates on the best performing model (in blue) compared to the naïve baseline (in orange).

External covariates have an overall positive impact on the forecasting performance. As for the solar energy production case, we notice a highly positive impact of the *future_weather* feature, which is expected as it contains windspeed information. The best performance is obtained by combining it with the time encoded features.

We also notice that the TiDE model is the only one outperforming the naïve baseline on this horizon and only when providing additional covariates, which shows the difficulty of forecasting wind power production.

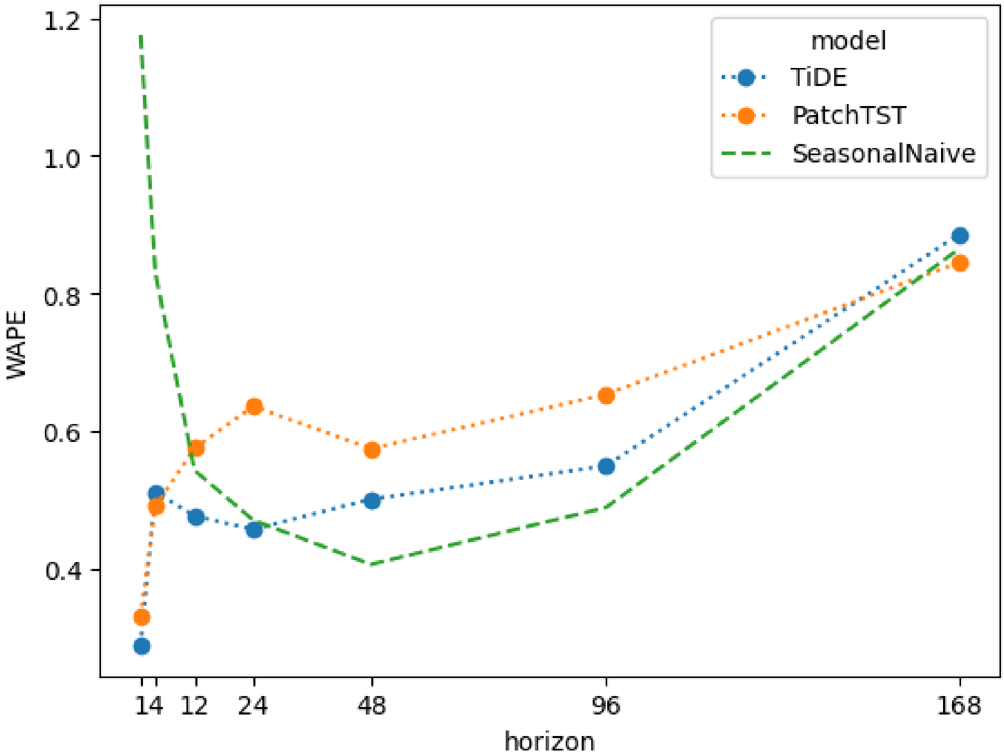

Figure 3-17: Impact of horizon on the best models' performance compared to the naïve baseline.

When observing the impact of horizon on the top models' performance, we notice that they mostly outperform the naïve baseline on short horizons (≤24h). Only the PatchTST model slightly outperforms the baseline with a horizon of 168h.

Table 3-4: Performance of evaluated models trained on a horizon of 24h with future_weather and time encoded features on test data.

Model	WAPE (%) ↓	MAE (kWh)	RMSE (kWh)
TiDE	44.35	5058	6344
Seasonal Naive	46.16	4969	6850
Weighted Ensemble	49.89	5922	7278
PatchTST	54.47	6648	8054
Chronos Bolt base (Zero Shot)	54.79	6395	8036
Chronos Bolt small (finetuned)	55.49	6673	8280
GBM	58.66	7338	8721
DeepAR	59.24	7053	8689
TFT	62.50	7749	9360

3.11 Network price forecasting

This case addresses the short-term forecasting of network and market prices, which will be a key enabler for planning and optimization algorithms across the U2Demo framework. Accurate price forecasts allow communities to make informed decisions not only for CO_2 reduction or self-consumption, but also for economic efficiency, flexibility scheduling, and peer-to-peer trading strategies.

3.11.1 Dataset

The dataset used for network price forecasting is the European Wholesale Electricity Price Data sourced from ENTSO-e, EMR and SEMOpx and cleaned by Ember [15]. It contains hourly wholesale electricity prices in €/MWh for 31 European countries spanning from 2015 to the present.

For this study, we focus solely on France market prices.

The data has already been cleaned by interpolating missing values from nearby ones.

3.11.2 Data treatment

There are a few missing values remaining in the dataset that we fill in with interpolation.

Encoded time features and lagged target values are added as exogenous covariates.

Data is split using sliding windows, we set aside the 5 last weeks of data from each site, 4 weeks of as validation data and the last week as test data to avoid any data leak. The rest of the data is used as training data for the models.

3.11.3 Method comparison

In this section, we examine the impact of different machine learning approaches on energy price forecasting performance in the French market. We compare a range of classical ML algorithms and deep learning architectures to assess their ability to capture nonlinear dynamics and temporal dependencies in price formation.

We evaluate the contribution of engineered features such as time encoded and lagged target features. We also measure how the models' performance evolve when increasing the forecasting horizon.

3.11.4 Results

Table 3-5: Performance of evaluated models on a horizon of 24h with lags and time encoded features on test data

Model	WAPE (%) ↓	MAE (€/MWh)	RMSE (€/MWh)
DeepAR	58.09	14.06	17.62
Weighted Ensemble	61.06	13.68	17.15
Chronos Bolt base (Zero Shot)	64.95	14.20	18.18
Chronos Bolt small (finetuned)	69.20	14.48	18.42
PatchTST	76.83	17.58	21.69
Seasonal Naive	78.60	16.81	22.19
TiDE	78.68	18.25	22.03
GBM	82.82	16.27	20.26
TFT	90.35	17.11	20.80

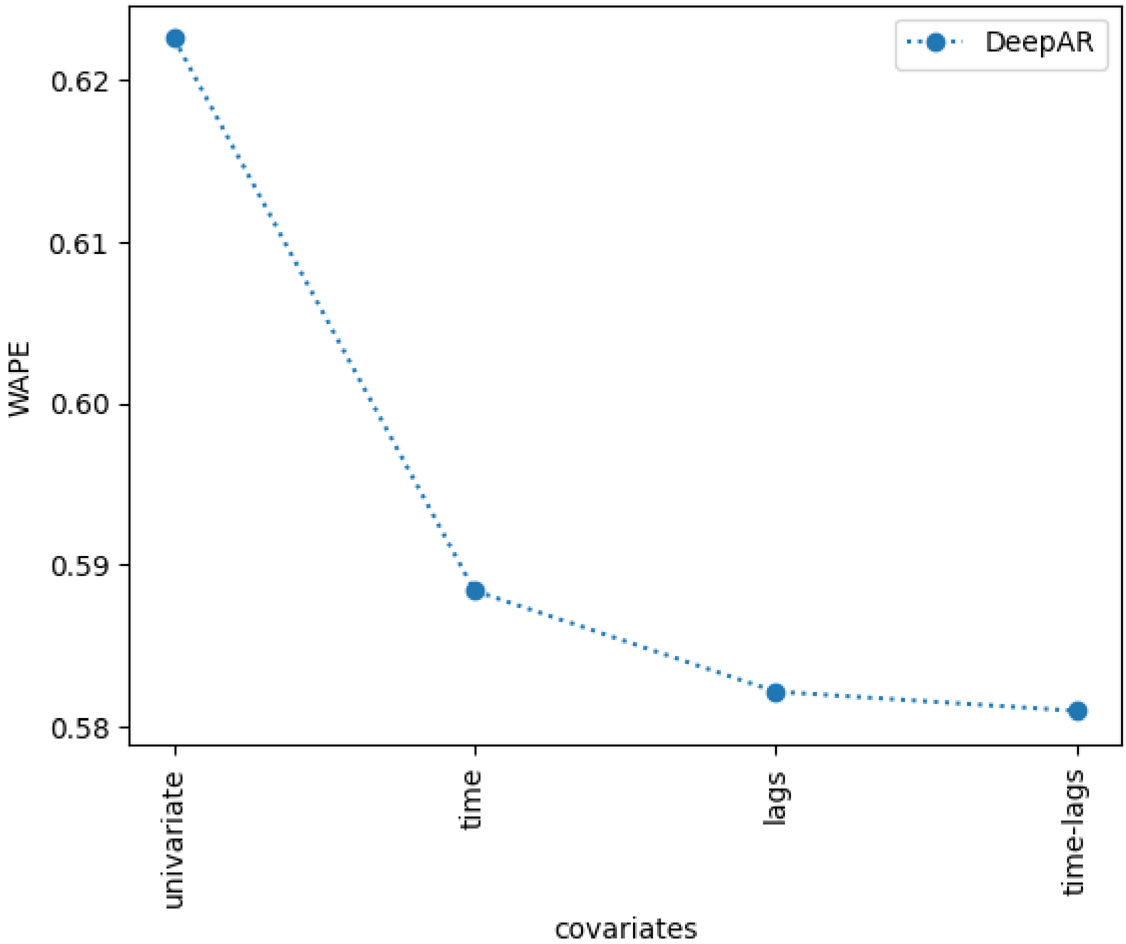

Figure 3-18: Impact of exogenous covariates on the best performing model.

We observe that time-encoded and lags covariates improve the DeepAR model performance, going from 62.26% in the univariate case down to 58.09% WAPE with both covariates (**Figure *3-18***).

On the influence of horizon on forecast performance, we note that the smallest error rates are achieved at 24h but also at 168h horizon while increasing in between (**Figure 3-19**).

Overall, the best performing model is DeepAR achieving 58.09% WAPE for an horizon of 24h corresponding to 14.06€/MWh MAE (**Table 3-5**).

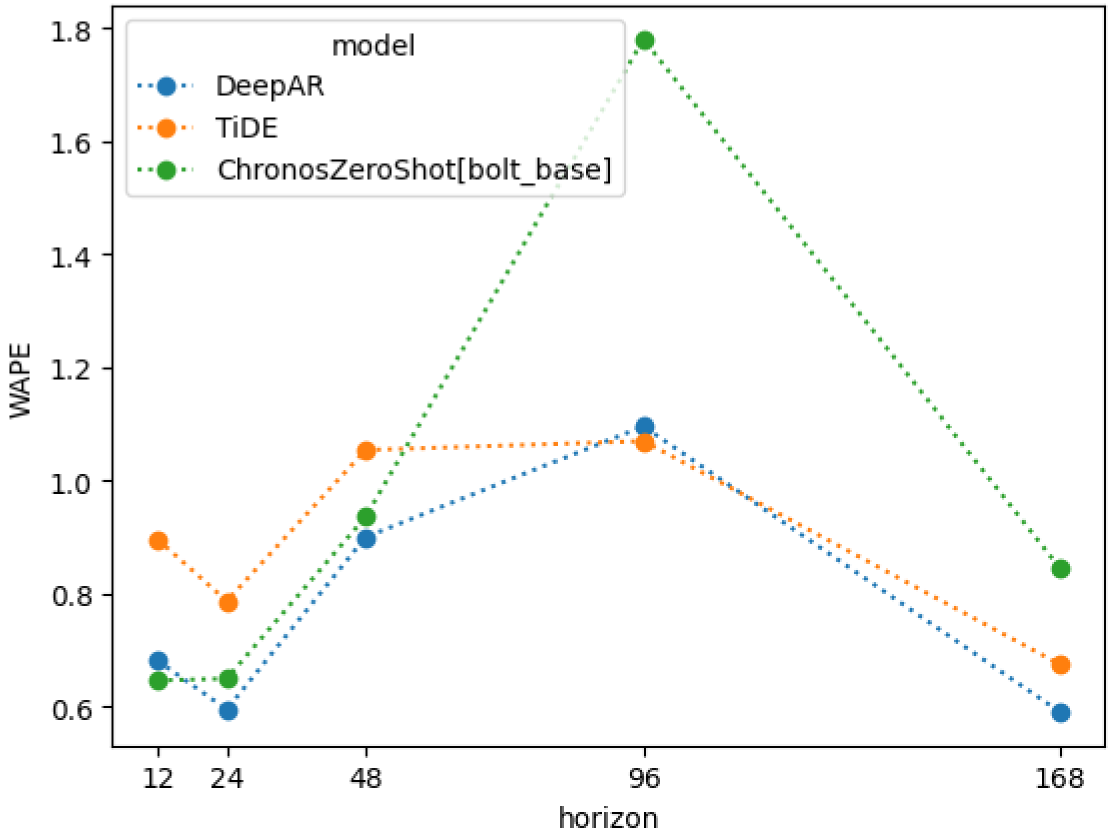

Figure 3-19: Impact of horizon on some of the best models' performance.

While the results are promising, energy market prices remain highly variable, making precise forecasting challenging. Therefore, using percentile forecasts rather than a single predicted value can be particularly valuable in this context (Figure 3-20).

Additionally, discrepancies between electricity market prices across different European countries must be considered in T4.2 for the various pilot sites.

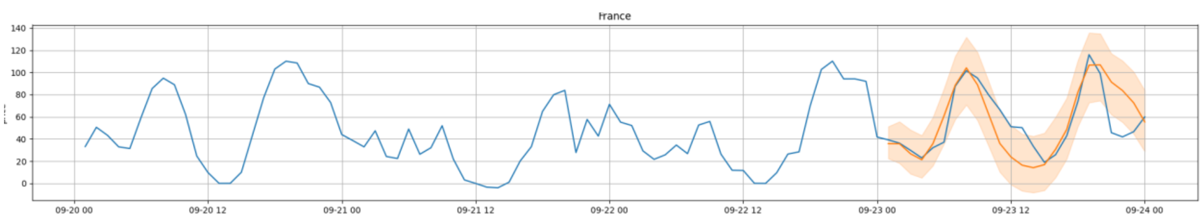

Figure 3-20: Example inference with a horizon of 24h on test data with the best DeepAR model trained with lagged target and time-encoded features. In blue is the observed data. The orange line is the mean forecasted value and the orange area represents the 10th-90th percentile interval.

3.12 General conclusion and recommendations

Across the scenarios studied we achieved solid predictive performance using a common, multi-horizon methodology and a consistent evaluation design. These results confirm the feasibility of deploying a forecasting layer that can support decision making for active consumers and energy communities in U2Demo. At the same time, it is crucial to stress that forecast accuracy is bounded by the *quantity* and *quality* of input data, coverage, resolution, label consistency, and the availability of truly forecast exogenous drivers (e.g., weather, market signals) all materially influence outcomes.

Several pilots may initially provide limited history and heterogeneous data streams. Under such conditions, models trained purely on local data can underperform or overfit. Therefore, we recommend that WP4.2 will explicitly compare local retraining versus transfer-learning strategies (e.g., pre-training on richer external datasets and fine-tuning on site data), selecting the approach that best balances accuracy, robustness, and operational effort. We also suggest introducing data-quality gates and drift monitoring to protect against silent degradation

The demo could face is come scenarios runtime, memory, or deployment limitations that restrict the use of heavier sequence-to-sequence models. We propose that WP4.2 baseline each deployment environment, evaluate and, where needed, introduce lightweight fallbacks (e.g., calibrated AR baselines), and package models for containerised execution with explicit resource envelopes. The aim should be to offer a graded portfolio, from fast, resilient baselines to richer models where infrastructure permits.

Forecasts in D2.2 were not yet integrated with the decision-support layer. We recommend that WP4.2 (i) align forecast horizons, update cadences, and latency requirements; (ii) define and standardize interfaces and data schemas for hand-off (including quantiles, confidence intervals, and scenario bundles); and (iii) co-tune objectives so that forecasts are fully compatible with scheduling, trading, and flexibility modules. This will ensure forecasting remains a key enabler for strategy and flexibility planning, as envisioned in this work package.

While the explored forecasting algorithms already compute diagnostics and summary statistics, we have not yet exploited retro-feedback from the data-analytics layer. We encouraged that in WP4.2 studies how feeding these analytics back into the feature space (or into model priors) impacts accuracy and stability.

It will be valuable to explore in WP4.2 weather forecasting certain flexibility scenarios is feasible and useful. For example, predicting conditions that trigger flexibility needs or opportunities. Even though the decision-support algorithms will ultimately select actions, assessing the potential of flexibility-oriented forecasts could enhance planning robustness and inform proactive strategies.

4 Data analytics methods

This section elaborates the role of data analytics in the development and assessment of methodologies that aid active customers and ECs (energy communities) in the definition of their strategies.

Energy communities promise lower bills, more renewable energy, and local benefits. Yet customers engage when they see personal relevance and agency: "my bill dropped, my rooftop PV is used locally, our community is fair". Data analytics links operational performance with such customer-facing value propositions. It quantifies how local trading and flexibility shift consumption, reduce peaks, and cut emissions; and it powers products (dynamic tariffs, peer-to-peer (P2P) trading, and carbon-aware scheduling), that members can understand and control.

4.1 Role of Data Analytics

Data analytics methods allow an evaluation of the behaviour (translated in energy usage patterns, electricity production patterns, and storage patterns) by the active consumers in energy communities. By applying these both before and after different strategies are implemented, their impact can be evaluated. This evaluation can be used to improve the efficiency and effectiveness of the strategies.

For this, the U2Demo project considers technical, economic, and social factors driving the data analytics.

In general, data analytics are categorized in three different categories.

- *Descriptive analytics* aim to understand what happened in the past, by analysing historical data and contextual information.
- *Predictive analytics* make forecasts on what will happen in the future.
- *Prescriptive analytics* recommend what can be done in order to reach certain goals or targets.

These data analytics can be translated into metrics or indicators (economic, technical and/or social), and they can be assessed at an individual or energy-community level, as elaborated in the following sub-sections.

Eventually, such data analytics can inform and engage the customers in the ECs.

Based on the metrics, descriptive analytics can be translated into dashboards for the individuals and the community. Examples include:

- Member app widgets: 'My savings this month' (EUR and %), 'My local energy share' (based on self-consumption and self-sufficiency), 'CO2 avoided' (kg), 'My P2P trades' (kWh, # counterparties), 'My peak contribution', etc.
- Community board: total savings, fairness indices, total local share, peak reduction, total CO2 avoided; trend vs baseline; collective goals and progress.

If the energy data is shown in combination with exogeneous factors (weather, market prices, community behaviour), trends can be highlighted.

The *predictive analytics* can expose the forecasts in apps such as 'tomorrow's plan', that can show expected member and community-level metrics.

The *prescriptive analytics* will provide decision support and show the effects of different optimisation scenarios or operational options, in terms of economic, technical or social metrics.

The following sub-sections elaborate these indicators, introducing a section-specific notation.

4.2 Technical Factors

Technical factors mostly relate to the electric quantities or ratios within the individual customer's premises (also called *members* of the EC), within the community as a whole, or of the grid to which they are connected.

4.2.1 Technical Factors at EC Level

At the level of the energy community, the following factors are the most important ones: SCR, SSR and peak active power.

4.2.1.1. The Self-Consumption Ratio of the EC

The self-consumption ratio measures how much of the electricity produced within an EC is consumed locally rather than exported to the main grid.
It is defined over a certain time period (e.g. day, month, year) and is expressed as a dimensionless quantity.

$$SCR = \frac{E_{self\_cons}}{E_{prod}} = \frac{E_{prod} - E_{exp}}{E_{prod}} \tag{4-1}$$

where

E_{self_cons}: Electrical energy ([kWh]) "self-consumed" in the community, by all members and shared assets, during the given time period.

E_{prod}: Electrical energy ([kWh]) generated ("produced") in the community, by all members and shared assets, during the given time period.

E_{exp}: Electrical energy ([kWh]) injected ("exported") into the main grid by the community, during the given time period.

The SCR ranges from 0 to 1, and can be interpreted as follows.

- SCR = 1 (100%) → All locally generated electricity is consumed within the community (ideal case).
- SCR close to 0 → Most of the generated energy is exported to the grid, meaning low local utilization.

A higher SCR means more community autonomy (less export) and better use of local renewables. It is often linked to higher savings because local consumption avoids retail tariffs and reduces grid fees. Additionally, it has a positive grid impact as higher self-consumption reduces stress on the distribution grid and lowers reverse power flows.

4.2.1.2. The Self-Sufficiency Ratio of the EC

The self-sufficiency ratio measures how much of the electricity consumed within an energy community is generated locally rather than imported from the main grid.
It is defined over a certain time period (e.g. day, month, year) and is expressed as a dimensionless quantity.

$$SSR = \frac{E_{self\_prod}}{E_{cons}} = \frac{E_{cons} - E_{imp}}{E_{cons}} \quad (4\text{-}2)$$

where

E_{self_prod}: Electrical energy ([kWh]) “self-produced” in the community, by all members and shared assets, during the given time period.

E_{cons}: Electrical energy ([kWh]) used (“consumed”) in the community, by all members and shared assets, during the given time period.

E_{imp}: Electrical energy ([kWh]) taken off (“imported”) from the main grid by the community, during the given time period.

The SSR ranges from 0 to 1, and can be interpreted as follows.

- SSR = 1 (100%) → All consumed electricity is locally generated within the community (ideal case).
- SSR close to 0 → Most of the consumed energy is imported to the grid.

A higher SSR means more community autonomy (less import) and better use of local renewables. It is often linked to higher savings because of reduced imports.

4.2.1.3. The Peak Active Power at EC level

The peak active power measures the maximal net power flow (offtake or injection) in the whole energy community in kW.

While it represents an instantaneous power value [kW], it is typically calculated for each time interval (e.g., 15 minutes, 1 hour) by dividing the energy used during that period [kWh] by the length of the interval [h]. The peak offtake or injection then corresponds to the maximum value of this instantaneous power observed over a defined time range.

$$peak_{EC}^{off} = \max_{t \in T} \sum_{i \in I} \left((E_{i,t}^{off} - E_{i,t}^{inj}) / \Delta t \right) \quad (4\text{-}3)$$

$$peak_{EC}^{inj} = \max_{t \in T} \sum_{i \in I} \left((E_{i,t}^{inj} - E_{i,t}^{off}) / \Delta t \right) \quad (4\text{-}4)$$

where

$E_{i,t}^{off}$: Electricity offtake (import) of member i at time instance t [kWh].

$E_{i,t}^{inj}$: Electricity injection (export) of member i at time instance t [kWh].

T : set of time intervals.

I : set of community members.

Δt : duration of time interval [h]

The peak active power is related to the grid connection contract of the EC, which specifies its maximum allowed peak load. Higer peaks stress the distributions networks more and can trigger thermal overloads. They are often linked to higher grid tariffs.

4.2.1.4. Grid Import

The grid import is computed based on the total energy imported from the grid (expressed in kWh):

$$\sum_{i \in I} E_i^{imp} \tag{4-5}$$

E_i^{imp} is the grid import energy for each of the EC members and I is the set of community members.

4.2.1.5. Battery Longevity

Battery longevity is computed based on number of battery cycles (expressed in battery cycles):

$$\sum_{i \in I} \sum_{b \in B} \sum_{t \in T} \frac{E_{i,b,t}^{chBESS} + E_{i,b,t}^{dchBESS}}{E_{i,b}^{maxBESS}} \tag{4-6}$$

where:

$E_{i,b,t}^{chBESS}$ is the energy charged by BESS units.

$E_{i,b,t}^{dchBESS}$ is the energy discharged by BESS units

$E_{i,b}^{maxBESS}$ is the maximum battery capacity from BESS units.

4.2.1.6. Demand Reduction

The comfort level is computed based on the total curtailed, reduced and unmet energy demand (expressed in kWh):

$$\sum_{i \in I} \sum_{\mathrm{l} \in \mathrm{L}} E_{i,l}^{loadRED} + E_{i,l}^{loadCUT} + E_{i,l}^{loadENS} \tag{4-7}$$

$E_{i,l,t}^{loadRED}$ represents voluntary demand reduction relative to the baseline schedule at time t for each load $l \in L$,

$E_{i,l,t}^{loadCUT}$ represents voluntary curtailment at time t for each load $l \in L$,

$E_{i,l,t}^{loadENS}$ represents energy not supplied by involuntary curtailment at time t for each load $l \in L$. L is the set of loads, while I represents the set of EC members.

4.2.2 Technical Factors at Individual Level

At the individual level, several static and dynamic metrics characterise the technical factors.

4.2.2.1. Static Individual Characteristics

A number of electric characteristics describe the individual capacities of community member i, relating to their grid connection, as well as individual assets. These are the most important ones.

Grid connection description:
- Single phase or three-phase AC connection: 230V, 3N400V, 3N230V, 3,230V, …
- Maximum current for injection or take-off: [A]
- Rated apparent power [kVA]
- Tolerable power factor (*cos phi*)

From these grid connection characteristics, one can derive the **maximum active power** to and from member i: $P_i^{max,offtake}$ [kW] and $P_i^{max,inj}$ [kW].

Installation description:
- Installed solar photovoltaic peak power (measured at the grid-side of the invertor) [kW]: P_i^{PV}.
- Installed battery power for charging and discharging (measured at the grid-side of the invertor) [kW]: $P_i^{batt,chg}$ and $P_i^{batt,dchg}$.
- Total battery capacity [kWh]: E_i^{batt} .
- Battery charging and discharging efficiency: $\eta_i^{batt,chg}$ and $\eta_i^{batt,dchg}$.

From these installation characteristics, one can derive the **peak amount of renewable power**, as well as different characteristics related to **flexibility**.

4.2.2.2. Dynamic Individual Characteristics

In contrast to the static characteristics, that don't change over time, the dynamic characteristics are key input to data analytics applications, as they reflect the actual varying status of quantities of the EC members.

4.2.2.3. The Individual Self-Consumption Ratio

The self-consumption ratio measures how much of the electricity produced within an individual member i of the energy community is consumed locally rather than exported to the main grid. It is defined over a certain time period (e.g. day, month, year) and expressed as a dimensionless quantity.

$$SCR_i = \frac{E_{self\_cons,i}}{E_{prod,i}} = \frac{E_{prod,i} - E_{exp,i}}{E_{prod,i}} \tag{4-8}$$

where

$E_{self\_cons,i}$: Electrical energy ([kWh]) “self-consumed” in the community, by member i, during the given time period.

$E_{prod,i}$: Electrical energy ([kWh]) generated (“produced”) by member i within the community during the given time period.

$E_{exp,i}$: Electrical energy ([kWh]) injected (“exported”) by member i during the given time period (this may include energy injected into the main grid or supplied to other peers/community members through local energy sharing mechanisms).

4.2.2.4. The Individual Self-Sufficiency Ratio

The self-sufficiency ratio measures how much of the electricity consumed by member i generated locally by member i rather than imported from the main grid.
It is defined over a certain time period (e.g. day, month, year), and expressed as a dimensionless quantity.

$$SSR_i = \frac{E_{self\_prod,i}}{E_{cons,i}} = \frac{E_{cons,i} - E_{imp,i}}{E_{cons,i}} \qquad (4\text{-}9)$$

where

$E_{self\_prod,i}$: Electrical energy ([kWh]) “self-produced” by member i, during the given time period.

$E_{cons,i}$: Electrical energy ([kWh]) used (“consumed”) by member i, during the given time period.

E_{imp}: Electrical energy ([kWh]) taken off (“imported”) from the main grid by member i (this may include energy imported from the main grid or supplied from other peers/community members through local energy sharing mechanisms), during the given time period.

4.2.2.5. The Peak Active Power at Individual Level

The peak active power measures the maximal net power flow (offtake or injection) in kW of member i.

While it represents an instantaneous power value [kW], it is typically measured as the energy [kWh] used during a certain time period (e.g. 15 minutes, 1 hour), divided by the length of the period [h].

$$peak_i^{off} = \max_{t \in T}\left((E_{i,t}^{off} - E_{i,t}^{inj})/\Delta t\right) \qquad (4\text{-}10)$$

$$peak_i^{inj} = \max_{t \in T}\left((E_{i,t}^{inj} - E_{i,t}^{off})/\Delta t\right) \qquad (4\text{-}11)$$

where

$E_{i,t}^{off}$: Electricity offtake (import) of member i at time instance *t* [kWh].

$E_{i,t}^{inj}$: Electricity injection (export) of member i at time instance *t* [kWh].

T : set of time intervals.

Δt : duration of time interval [h]

4.2.2.6. The Individual Member Interacting with the EC

Several characteristics explain the activity of the individual member i withing the EC.

First, we define the **total amount of electrical energy that a member *i* sells or shares on the P2P market** over a time horizon *T* (e.g. day, week, month). It is expressed in [kWh].

$$E_i^{p2p\_sell} = \sum_{t \in T} E_{i,t}^{p2p\_sell} \tag{4-12}$$

where

$E_{i,t}^{p2p\_sell}$: amount of electrical energy that a member i sells on the P2P market during a period Δt, at time *t* [kWh].

T : set of time intervals.

Similarly, we define the t**otal amount of electrical energy that a member *i* buys or receives from others on the P2P market** over a time horizon *T* (e.g. day, week, month). It is expressed in [kWh].

$$E_i^{p2p\_buy} = \sum_{t \in T} E_{i,t}^{p2p\_buy} \tag{4-13}$$

where

$E_{i,t}^{p2p\_buy}$: amount of electrical energy that a member i buys on the P2P market during a period Δt , at time *t* [kWh].

T : set of time intervals.

These two characteristics help defining the 'success' of the P2P market. If we compare it with the overall electricity bought/sold, we get its importance relative to the retail market with the suppliers.

As such we can define the **percentage of electricity offtake that comes from the P2P market** for each member [%]:

$$E_{relative,i}^{p2p\_buy} = \frac{\sum_{t \in T} E_{i,t}^{p2p\_buy}}{\sum_{t \in T} E_{i,t}^{off}} \tag{4-14}$$

As well as the **percentage of electricity offtake that comes from the retailer** [%]:

$$E_{relative,i}^{retailer\_buy} = 1 - \frac{\sum_{t\in T} E_{i,t}^{p2p\_buy}}{\sum_{t\in T} E_{i,t}^{off}} \quad (4\text{-}15)$$

4.2.3 Technical Factors at Grid Level

The energy community is connected to the low voltage distribution grid, on which certain power quality requirements are imposed. The energy community and its individual members should not worsen grid problems (such as congestion) or can even help solving grid problems if the distribution systems operator can participate or intervene on the P2P market.

As they are not the main focus of this project, they are only mentioned briefly here.

IEC standard EN 50160, which is a harmonized standard in the European Union, entitled: "Voltage characteristics of electricity supplied by public distribution systems" provides the most important specifications of the grid power quality.

As such, the ***rms* value of the voltage** needs to be between 0.9 and 1.1 per unit, for 95% of time. The voltage can drop even below 85%, but only for less than 1 minute and at most 1000 times per year.

Hence for data analytics purposes, this standard allows defining **maximum and minimum voltage levels**, as well as the frequency of overvoltage/undervoltage events.

A different aspect of congestion is thermal overloading of lines and assets such as transformers. This can be translated into feeder/branch use characteristics, such as the ratio of the actual load power of a branch to its rated power, as well as the frequency of such overload events.

Other grid aspects deal with unbalance between the different phases of the 3-phase system. From EN 50160, the **Voltage Unbalance Factor** (VUF) should remain below 2%, for 95% of time. Finally, harmonics distortion shall be limited, violations of these limits can be used in data analytics applications as well.

4.3 Economic Factors

Besides the technical factors, it is important to consider economic factors as well. They can be identified at the individual level, as well as at the EC level.

4.3.1 Individual Cost Savings (absolute value / percentage)

Individual cost savings in an energy community quantify how much a member reduces their electricity cost by participating in local trading (or other community mechanisms) compared to a baseline scenario without those mechanisms. Hence, for member i, the **individual cost savings** CSi are expressed in [€] and defined over a predefined time period (day, week, month).

$$CS_i = C_i^{base} - C_i^{P2P} \quad (4\text{-}16)$$

where

C_i^{base}: baseline cost without P2P trading over the time period.

This baseline cost can be calculated based on offtake and injection information of each individual member.

$$C_i^{base} = \sum_{t\in T}\left(E_{i,t}^{off} \cdot \lambda_t^{off} - E_{i,t}^{inj} \cdot \lambda_t^{inj}\right) \quad (4\text{-}17)$$

where

$E_{i,t}^{off}$, $E_{i,t}^{inj}$: Offtake/injection quantity of consumer i in kWh in each time step.

λ_t^{off}, λ_t^{inj}: Retail offtake/injection price in €/kWh in each time step.

T : set of time intervals.

C_i^{P2P}: the cost when peer-2-peer trading or energy sharing is active, over the same time period.

From this, one can also derive the relative cost savings, or **individual cost saving index in percentage**:

$$CSI_i = \frac{C_i^{base} - C_i^{P2P}}{C_i^{base}} \times 100\%\% \quad (4\text{-}18)$$

As EC members want to see tangible financial benefits; this metric is the most direct way to communicate value.

4.3.2 EC-level Cost Savings (absolute value / percentage)

By adding/aggregating the information over all members, one can define the EC-level cost savings as the aggregated financial benefit for all members compared to the baseline scenario without local trading or community mechanisms.

Total cost savings (CS) are expressed in [€] and defined over a predefined time period.

$$CS = \sum_{i\in I} CS_i = \sum_{i\in I} C_i^{base} - \sum_{i\in I} C_i^{p2p} \quad (4\text{-}19)$$

where

I : set of community members.

Then, the total **cost saving index in percentage** describes the percentage reduction in total community energy costs relative to the baseline. It is defined as

$$CSI = \frac{\sum_{i\in I} C_i^{base} - \sum_{i\in I} C_i^{p2p}}{\sum_{i\in I} C_i^{base}} \text{x}100\% \quad (4\text{-}20)$$

These values identify the overall savings of the entire community, and can be useful for engaging customers *("Together we saved €X this month")* or compare performance across months or other communities.

4.3.3 Total Operational Cost

The total operation cost is computed based on the total energy invoice and operational costs of equipment (expressed in €):

$$\sum_{i\in I}\sum_{t\in T} E_{i,t}^{imp}\cdot\lambda_{i,t}^{imp} - E_{i,t}^{exp}\cdot\lambda_{i,t}^{exp} + \sum_{i\in I}\sum_{b\in B}\sum_{t\in T} E_{i,b,t}^{chBESS}\cdot\lambda_{i,b,t}^{chBESS} + E_{i,b,t}^{dchBESS}\cdot\lambda_{i,b,t}^{dchBESS} + \sum_{i\in I}\sum_{ev\in EV}\sum_{t\in T} E_{i,ev,t}^{chEV}\cdot\lambda_{i,ev,t}^{chEV} + E_{i,ev,t}^{dchEV}\cdot\lambda_{i,ev,t}^{dchEV} + \sum_{i\in I}\sum_{l\in L}\sum_{t\in T} E_{i,l,t}^{loadRED}\cdot\lambda_{i,l,t}^{loadRED} + E_{i,l,t}^{loadCUT}\cdot\lambda_{i,l,t}^{loadCUT} + \sum_{i\in I}\sum_{g\in G}\sum_{t\in T} E_{i,g,t}^{genAct}\cdot\lambda_{i,g,t}^{genAct} + E_{i,g,t}^{genExc}\cdot\lambda_{i,g,t}^{genExc} \quad (4\text{-}21)$$

$E_{i,t}^{imp}$ and $E_{i,t}^{exp}$ are the grid import and export energy at time t respectively, with $\lambda_{i,t}^{imp}$ and $\lambda_{i,t}^{exp}$ representing their corresponding cost for each member $i \in I$.

$E_{i,g,t}^{genAct}$ is the energy used by the community at time t, with $\lambda_{i,g,t}^{genAct}$ representing its operating or opportunity cost for each generator $g \in G$, for each member $i \in I$.

$E_{i,g,t}^{genExc}$ is the surplus energy that is not utilized internally (spilled or curtailed), and $\lambda_{i,g,t}^{genExc}$ is the corresponding penalty or opportunity cost for each generator $g \in G$, for each member $i \in I$.

$E_{i,l,t}^{loadRED}$ represents voluntary demand reduction relative to the baseline schedule at time t, and $\lambda_{i,l,t}^{loadRED}$ is the cost associated with that reduction for each load $l \in L$, for each member $i \in I$.

$E_{i,l,t}^{loadCUT}$ represents voluntary curtailment at time t, and $\lambda_{i,l,t}^{loadRED}$ is the corresponding cost for each load $l \in L$, for each member $i \in I$.

$E_{i,b,t}^{chBESS}$ and $E_{i,b,t}^{dchBESS}$ are the energy charged and discharged, with $\lambda_{i,b,t}^{chBESS}$ and $\lambda_{i,b,t}^{dchBESS}$ their respective marginal costs (e.g., cycling, efficiency losses) for each BESS unit $b \in B$, for each member $i \in I$.

$E_{i,ev,t}^{chEV}$ and $E_{i,ev,t}^{dchEV}$ are the energy charged and discharged with $\lambda_{i,ev,t}^{chEV}$ and $\lambda_{i,ev,t}^{dchEV}$ being the corresponding costs for each electric vehicle $ev \in EV$, for each member $i \in I$.

4.3.4 Total energy invoice

The total energy invoice is computed based on the total cost of importing energy from the grid and the total remuneration of exporting energy to the grid (expressed in €):

$$\sum_{i\in I}\sum_{t\in T}\left(E_{i,t}^{imp}\cdot\lambda_{i,t}^{imp} - E_{i,t}^{exp}\cdot\lambda_{i,t}^{exp}\right) \quad (4\text{-}22)$$

4.3.5 Maximum Potential Saving Index

One can also define the **maximum potential saving index MPSI** for the entire community, a metric that measures how close the actual cost savings of an energy community are to the theoretical maximum possible savings if all locally generated electricity were consumed within the community.

$$MPSI = \frac{C_{base} - C_{P2P}}{C_{base} - C_{min}} \tag{4-23}$$

where

C_{min}: minimum cost that the whole P2P community would pay over a predefined time period – assuming that all generated electricity can be locally consumed.

$$C_{min} = \sum_{t \in T} \left[\lambda_t^{off} \cdot \max(0, \Lambda_t) - \lambda_t^{inj} \cdot \max(0, -\Lambda_t) \right] \tag{4-24}$$

where

$$\Lambda_t = \sum_{i \in I} \left(E_{i,t}^{off} - E_{i,t}^{inj} \right) \tag{4-25}$$

This *MPSI* shows how efficiently the community uses its local generation and can be used as an engagement tool (e.g. “We’re at 65% of our maximum potential—help us reach 80%!”), as well as for optimization purposes, guiding strategies for storage, demand response, and P2P pricing.

4.4 Social and Environmental Factors

Finally, it is important to consider social and environmental factors as well. Most of these can be defined at the community level, but one can also derive some individual level.

These factors capture the fairness, inclusiveness, and environmental impact of an energy community. These dimensions influence trust, participation, and long-term engagement. While technical and economic metrics show efficiency and cost benefits, social indicators ensure that these benefits are shared equitably and align with community values.

4.4.1 Jain’s Fairness Index

Jain's fairness index is a quantitative metric used to evaluate the fairness of resource allocation among multiple users in a system. It is used in different domains. The index ranges from $\frac{1}{n}$ (the least fair allocation, where only one user gets all resources) to 1 (perfect fairness, where all users get equal shares). Jain's index is independent of the scale and metric—meaning its value doesn't change if all allocations are multiplied by a constant.

In energy communities, Jain's Fairness Index can be based on the percentage cost savings CSIi or the absolute cost savings CSi as follows

$$FI = \frac{\left(\sum_{i \in I} CSI_i\right)^2}{|I| \cdot \sum_{i \in I} CSI_i^2} = \frac{\left(\sum_{i \in I} CS_i\right)^2}{|I| \cdot \sum_{i \in I} CS_i^2} \quad (4\text{-}26)$$

where

$|I|$ is the number of members in the community set I

while the other items have been defined above.

As an interpretation, if FI → 1, one reaches full fairness (everyone benefits equally), while lower FI means higher disparity in benefits.

It captures the 'fair share' perspective: does every member has an equal part? As such, come consider it is more an equity index than a fairness index, as it does ignore the individual contributions.

4.4.2 Gini Coefficient

The Gini coefficient or Gini fairness index (GI) is defined differently, where lower values mean more equality or fairness. (Hence, opposite to *FI*.)

$$GI = \frac{\sum_{i \in I} \sum_{j \in I} |CS_i - CS_j|}{2 \cdot |I| \cdot \sum_{i \in I} CS_i} \quad (4\text{-}27)$$

As an interpretation, if GI → 1, one reaches maximum inequality, while lower GI means higher equality.

It captures the 'gap between member' perspective: are all members considered appropriately?

4.4.3 MinMax Index

The MinMax Index is defined as the ratio of the smallest to the largest benefit. The closer it is to 1, the more fairly the benefits are shared. This index might be sensitive to outliers

$$MinMax = \frac{\min(CS_i)}{\max(CS_i)} \quad (4\text{-}28)$$

It captures the 'gap between extremes': how much better of is the best versus the worst?

4.4.4 Carbon Emissions Reduction

Finally, one can also quantify the **carbon emissions reduction for the community**, comparing the P2P case, to one without P2P trading as a baseline, where the electricity taken from the grid is considered with it (potentially time-varying) emission factor.

It is expressed as $kg_{CO2eq,}$ over a given time period.

$$\Delta CO2_{P2P} = \left(E_{net} - E_{net}^{base}\right) \cdot EF \quad (4\text{-}29)$$

where

E_{net}: Total net offtake electricity from the grid [kWh] during that time period.

E_{cons}^{base}: Total net offtake electricity from the grid [kWh] without P2P trading (=baseline).

EF: Emission factor [kg_{CO2eq}/kWh] based on the electricity generation mix in the grid.

This can also be considered at individual-level. The **Individual carbon emissions reduction** including individual self-consumption and collective self-consumption (P2P trading) over a period can be defined as follows.

$$\Delta CO2_i = \sum_{t\in T} \left(E_{i,t}^{gen} - E_{i,t}^{inj} + E_{i,t}^{ps}\right) \cdot EF \quad (4\text{-}30)$$

where

$E_{i,t}^{gen}$: Generated electricity of consumer i in kWh in each time step.

$E_{i,t}^{inj}$: Injection quantity of consumer i in kWh in each time step.

$E_{i,t}^{ps}$: Electricity sold on the P2P market by consumer i in kWh in each time step.

These carbon emissions reduction factors demonstrate the environmental impact to the community stakeholders.

4.4.5 Environmental Impact

The environmental impact is computed based on the total energy imported from the grid (expressed in kWh):

$$\sum_{i\in I} \sum_{t\in T} E_{i,t}^{imp} \cdot EF_t \quad (4\text{-}31)$$

4.5 Conclusion on Data Analytics

Data analytics plays a pivotal role in empowering energy communities (ECs) and their members to make informed, impactful decisions. By translating complex energy behaviors into actionable insights, analytics bridges the gap between technical performance and customer-facing value. Through descriptive, predictive, and prescriptive methods, it enables the evaluation and optimization of strategies that enhance local energy use, reduce costs, and promote sustainability.

At the heart of this approach are metrics that span technical, economic, social, and environmental dimensions. Technical indicators such as self-consumption and self-sufficiency ratios, peak power flows, and grid interaction parameters provide a detailed understanding of energy dynamics at both individual and community levels. Economic metrics quantify cost savings and efficiency, offering tangible incentives for participation. Social indicators like fairness indices ensure equitable benefit distribution, fostering trust and inclusiveness. Environmental metrics, particularly carbon reduction calculations, highlight the ecological benefits of local energy strategies.

These analytics not only support operational improvements but also serve as engagement tools. Dashboards and apps visualize savings, local energy shares, and environmental impact, making data accessible and meaningful to members. Predictive tools guide future planning, while prescriptive analytics offer decision support for optimization scenarios.

Ultimately, data analytics transforms energy communities from passive consumers into active participants. It enables transparency, drives performance, and aligns individual actions with collective goals. As ECs evolve, robust analytics will remain essential in ensuring their success, scalability, and sustainability.

5 Decision support methods

5.1 Main Objective

Members of an EC, although geographically close, may exhibit different consumption preferences. When EC managers have control over the EC member assets, these preferences should be included in the community optimal management, for the scheduling to be aligned with expectations of EC members. Having insights onto the expected scheduling is also beneficial for EC members to choose a specific or multiple consumption preferences – thus becoming a decision support method. The expected scheduling of each member will have a direct impact on the participation in energy sharing or peer-to-peer transactions and also on flexibility services.

The main objective of this chapter is to propose decision support methods for EC members to gain insights into the expected scheduling of flexible devices according to individual EC member consumption preferences. The insights are provided by the analysis of the expected scheduling according to KPIs defined in Section 4.

5.2 Use Case (UC)

The optimization is modelled using Python Energy Communities (PyECOM), a user-friendly, Python-based tool developed by INESC ID for the analysis, simulation, and optimisation of energy communities [16]. The simulations were carried out using the deterministic linear model implemented in PyECOM. The algorithm was implemented in Python using the Pyomo library and solved with the commercial solver CPLEX.

The use case (UC) considered in this chapter is a simulated UC of an EC comprising five prosumers. The community is equipped with five battery energy storage systems (BESS), as well as five electric vehicle supply equipment (EVSE) units with V2G technology, five electric vehicles (EV) that support V2G technology, and five photovoltaic (PV) generators, each of the members with one per location. The input data of consumption and production was obtained from a real dataset of five households on Madeira Island (Portugal) covering the second half of 2019. In addition, the technical constraints of the distribution network were incorporated into the analysis, including a local transformer with a nominal capacity of 50 kVA, subject to an operational limit of 80%. The EC's total load may be curtailed, reduced, or not supplied, but the sum of these components must not exceed 10% of the total load.

Figure 5-1 and Figure 5-2 present measured data from 6 November 2019 for the PV and load of each EC member. This data is considered as forecasted data for the proposed method. The profile is atypical, particularly at the early hours, with a peak at around 3 am, but since it the data is real, atypical data creates an opportunity for testing.

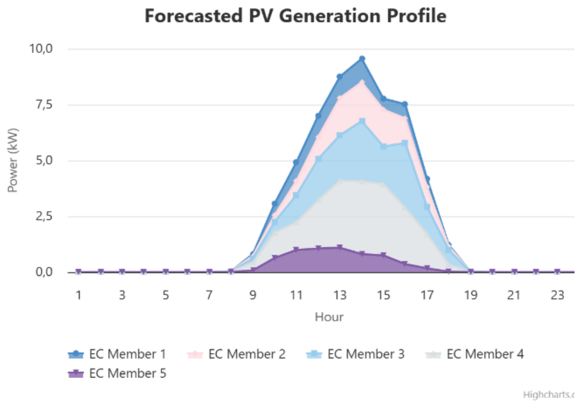

Figure 5-1: Forecast PV Generation Profile

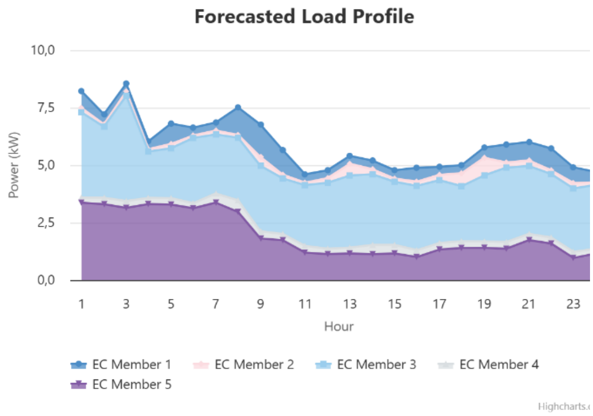

Figure 5-2: Forecasted Load Profile

Table 5-1 presents per-member information for the energy community, including each member's total load consumption, total generation, and total EV energy requirement.

Table 5-1: Information about the profile of the EC members

EC Member / Parameters	Member 1	Member 2	Member 3	Member 4	Member 5	Total
Total Load Power Consumption [kW]	14.22	5.71	68.63	7.16	47.25	142.97
Total Generator Power Production [kW]	5.94	9.29	14.87	18.79	5.94	54.83
Total EV Energy Required [kWh]	20.00	21.30	35.56	21.30	22.50	120.66

Table 5-2 presents information of the five EVs of the EC, including connection times (arrival and departure periods), charging prices that account for battery cycling, maximum energy capacity, and initial and target State-of-Charge (SoC).

Table 5-2: Information about the EC member's EVs.

Session ID	EV ID	EVSE ID	Max Capacity (kWh)	Initial SOC (%)	Target SOC (%)	Arrive Time Period	Departure Time Period
1	1	1	40.0	30	80	2	9
2	2	2	35.5	40	100	1	8
3	3	3	50.8	20	90	14	23
4	4	4	35.5	30	90	18	23
5	5	5	45.0	40	90	9	17

Figure 5-3 complement the information about the connection state of the five EVs of the EC.

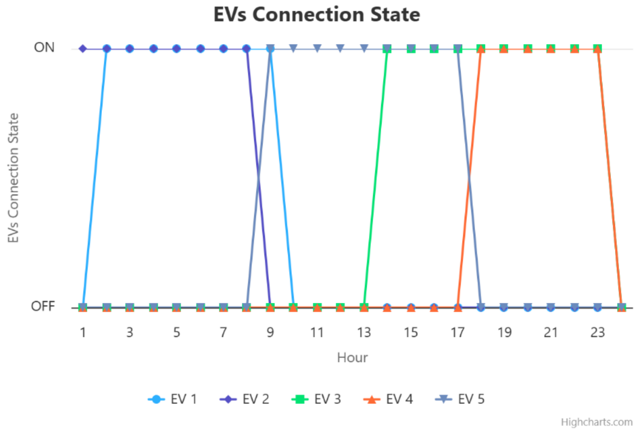

Figure 5-3: EVs connection state

Table 5-3 provides an overview of the technical parameters of the EC's BESSs.

Table 5-3: Information about EC's BESSs:

BESS ID	Max Capacity (kWh)	Initial SOC (%)	Target SOC (%)	Charge Efficiency (%)	Discharge Efficiency (%)
1	15	50	50	95	95
2	15	50	50	96	96
3	30	50	50	95	95
4	15	50	50	95	95
5	30	50	50	95	95

Table 5-4 summarizes the technical characteristics of the EC's EVSEs.

Table 5-4: Information about EC's EVSEs

EVSE ID	Charge Power Limit (kW)	Discharge Power Limit (kW)	Charge Efficiency (%)	Discharge Efficiency (%)
1	7.4	7.4	95	95
2	7.4	7.4	95	95
3	7.4	7.4	95	95
4	7.4	7.4	95	95
5	7.4	7.4	95	95

The emission factor for the energy imported from the grid was computed based on the reports published by the Portuguese Environment Agency [17] and by IPCC (Intergovernmental Panel on Climate Change) [18] together with the daily electricity balance data provided by REN (Redes Energéticas Nacionais) [19].

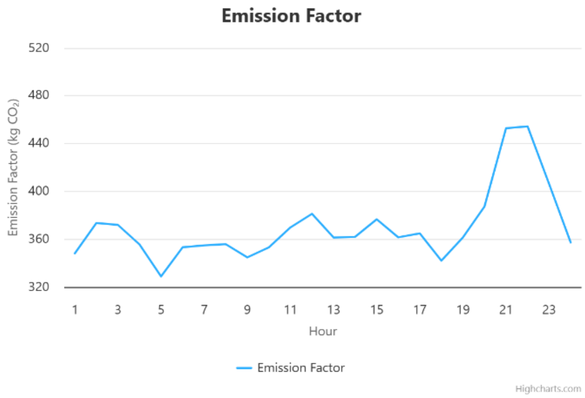

Figure 5-4: Emission Factor

5.3 Objective Functions and Multi-Objective Function

In this section, the objective functions (OFs) that translate the EC member consumption preferences are introduced together with the multi-objective function.

The OFs are expressed as follows:

- Baseline – EV charge as soon as possible and direct self-consumption (result expressed in kWh):

$$min\ F_1 = \sum_{ev \in EV} \sum_{t \in T} \frac{E_{ev}^{maxEV} - E_{ev,t}^{EV}}{\Delta t} + P_t^{imp} \qquad (5\text{-}1)$$

where E_{ev}^{maxEV} represents the EV battery maximum capacity, $E_{ev,t}^{EV}$ represents the EV current state of charge (SoC), P_t^{imp} represents the grid import power and Δt represents the optimisation timestep. In this section, it is considered that the timestep is one hour.

- Reduce operational cost (result expressed in €) – reduce the total cost considering energy invoice costs and the assets costs (e.g. charging and discharging costs for EVs and BESS, generator maintenance and load services activation costs):

$$\begin{aligned} min\ F_2 = &\sum_{t \in T} P_t^{imp} \cdot \lambda_t^{imp} - P_t^{exp} \cdot \lambda_t^{exp} + \sum_{b \in B} \sum_{t \in T} P_{b,t}^{chBESS} \cdot \lambda_{i,b,t}^{chBESS} + P_{b,t}^{dchBESS} \\ &\cdot \lambda_{b,t}^{dchBESS} + \sum_{ev \in EV} \sum_{t \in T} P_{ev,t}^{chEV} \cdot \lambda_{ev,t}^{chEV} + P_{ev,t}^{dchEV} \cdot \lambda_{ev,t}^{dchEV} \\ &+ \sum_{l \in L} \sum_{t \in T} P_{l,t}^{loadRED} \cdot \lambda_{l,t}^{loadRED} + P_{l,t}^{loadCUT} \cdot \lambda_{l,t}^{loadCUT} \\ &+ \sum_{g \in G} \sum_{t \in T} P_{g,t}^{genAct} \cdot \lambda_{g,t}^{genAct} + P_{g,t}^{genExc} \cdot \lambda_{g,t}^{genExc} \end{aligned} \qquad (5\text{-}2)$$

- Reduce energy invoice – minimize electricity costs (result expressed in €):

$$min\ F_3 = \sum_{t \in T} P_t^{imp} \cdot \lambda_t^{imp} - P_t^{exp} \cdot \lambda_t^{exp} \quad (5\text{-}3)$$

- Reduce grid import (result expressed in kWh):

$$min\ F_4 = \sum_{t \in T} P_t^{imp} \quad (5\text{-}4)$$

- Increase battery longevity – preserve battery's lifespan (result expressed in kWh):

$$min\ F_5 = \sum_{b \in B} \sum_{t \in T} P_{b,t}^{chBESS} \quad (5\text{-}5)$$

- Reduce environmental impact – reduce CO2 emissions (result expressed in kgCO2 kWh):

$$min\ F_6 = \sum_{t \in T} P_t^{imp} \cdot EF_t \quad (5\text{-}4)$$

- Increase comfort – Prioritize member routines: no delay or restriction on appliance operation (result expressed in kWh):

$$min\ F_7 = \sum_{l \in L} \sum_{t \in T} P_{l,t}^{loadCUT} + P_{l,t}^{loadRED} + P_{l,t}^{loadENS} \quad (5\text{-}7)$$

An EC member may also pursue multiple objectives simultaneously. Then, these OFs can be combined into a multi-objective function. In the proposed approach, this is done by assigning weights to the different OFs, thereby reflecting their relative importance. To this end, the multi-objective function can be formulated as follows:

$$min\ F = \sum_{i \in I} \sum_{n \in N} \alpha_{i,n} \cdot \beta_{i,n} \cdot F_{i,n} \quad (5\text{-}5)$$

where $i \in I$ refers to the EC members and $n \in N$ refers to the OFs. The weights $\alpha_n \in [0,1]$ set the contribution of each objective: when an objective is deactivated, $\alpha_n = 0$; when only one objective is considered, $\alpha_n = 1$ and the others are zero; in the multi-objective case, the weights satisfy $\sum_{n \in N} \alpha_n = 1$. The scaling factors $\beta_n > 0$ normalize the objectives, ensuring comparable magnitudes. **Table *5-5*** shows the OF correlation with β_n factor.

Table 5-5: OF correlation with β_n factor

OF	OF result unit	β_n	β_n Unit
Baseline	kWh	$\overline{C_{imp}}$	€/kWh
Reduce operational cost	€	1	–
Reduce energy invoice	€	1	–
Reduce Grid Import	kWh	$\frac{\overline{C_{imp}} + \overline{C_{exp}}}{2}$	€/kWh
Battery Longevity	kWh	$\overline{C_{imp}}$	€/kWh
Environmental Impact	Kg_{CO2} kWh	$\frac{\overline{C_{imp}}}{\overline{EF}}$	€/kWh · 1/(Kg_{CO2})
Comfort	kWh	$\overline{C_{imp}}$	€/kWh

In **Table 5-5**, $\overline{C_{imp}}$ represents the average import price, $\overline{C_{exp}}$ represents the average export price and $\overline{\mathrm{EF}}$ that represents the average emission factors.

5.4 Results – OFs considered individually

The main results obtained for the EC optimal scheduling are presented in the following subsections.

5.4.1 Baseline

Figure 5-5 shows the EC production profile under the Baseline OF. Generator power production is the PV generation utilized by the EC, either to meet internal demand or to export to the grid. BESS discharging power is the aggregate discharge from all community BESS. EV discharging power is the power delivered by the EV fleet. Load reduction denotes voluntary demand reductions relative to the original load. Load cut (curtailment) refers to voluntary service curtailment (not delivered). Load ENS represents involuntary load curtailment. Grid import power is the community's total power drawn from the grid, and the import price is the corresponding tariff at each time step.

Figure 5-6 shows the EC consumption profile under the Baseline OF. Load power consumption is the demand from EC members, after accounting for reductions and curtailments. BESS charging power is the aggregate charging across all community BESS, and EV charging power is the aggregate charging by the EV fleet. Generator power surplus denotes PV energy not utilized by the EC. Grid export power is the power delivered from the EC to the grid, and the grid export price is the corresponding export tariff at each time step.

Under the baseline OF, each EV initiates charging immediately upon connection, while the batteries operate in self-consumption mode. No smart scheduling strategy is applied, nor are energy tariffs taken into consideration. Although it aims to maximize self-consumption by minimizing grid imports, there's still grid import in almost every hour to meet the cars' demand, as can be seen in Figure 5-5. Furthermore, based on the results presented in Figure 5-6, there is no grid export power or V2X discharging. However, there is generator power surplus in hours of high PV production, with a peak of 1.38 kW at hour 13, that shows the inefficiency of the baseline OF to manage surplus from PV resources.

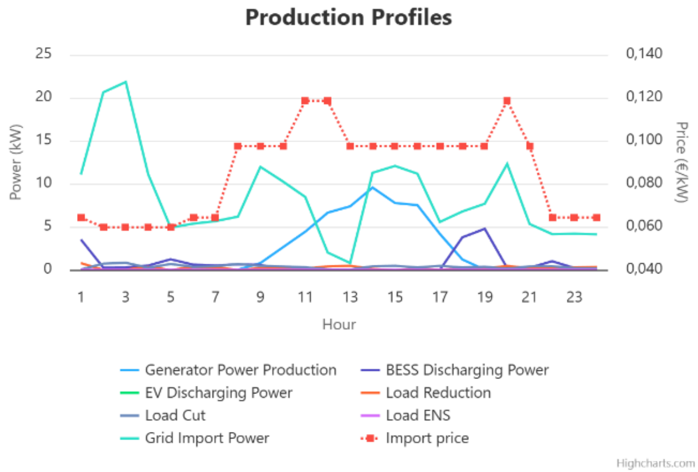

Figure 5-5: Baseline OF production profile

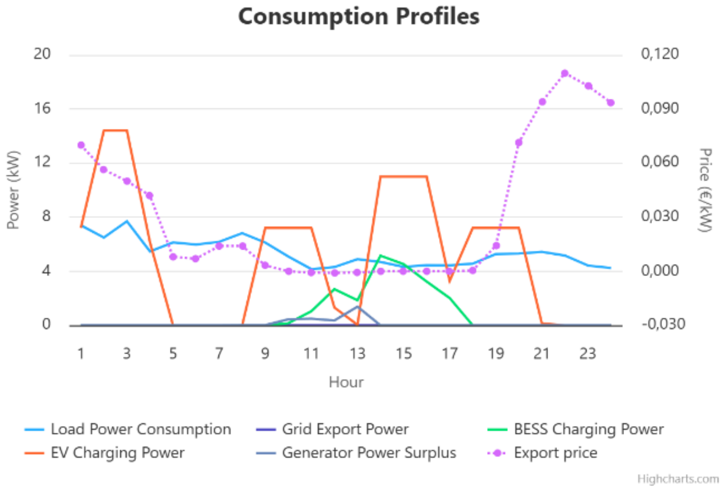

Figure 5-6: Baseline OF consumption profile

Figure 5-7 presents the total load that was curtailed, reduced, or not supplied during the optimization horizon. In total, demand decreased by 9.75% (13.93 kWh). This was achieved by 8.80 kWh of load cutting and 5.13 kWh of load reduction. Also, there was zero load not supplied.

Load Cut, Reduced and Not Supplied

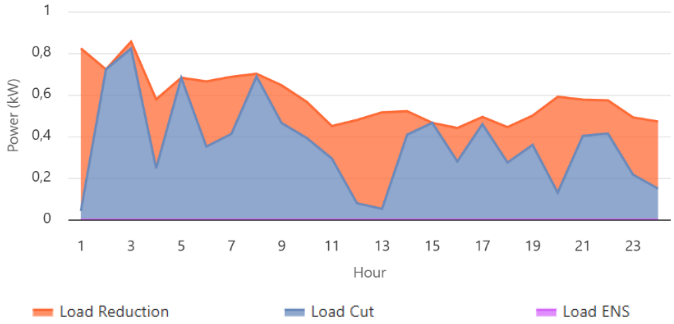

Figure 5-7: Baseline OF Load Cut, Reduced and Not Supplied

The BESS 1, 2 and 4 discharged to supply the load consumption and EV charging and charged during peak PV production hours (**Figure 5-8**). This behavior aligns with the import-minimization objective that is a component of the Baseline OF. All EVs reach their target SoC (**Figure 5-9**). In this figure, the value of the SoC arrives to zero because the EVs are not connected to the EC and the optimization don't have information about the real state of the SoC. At the end of the day, the total energy invoice of this approach is €17.11.

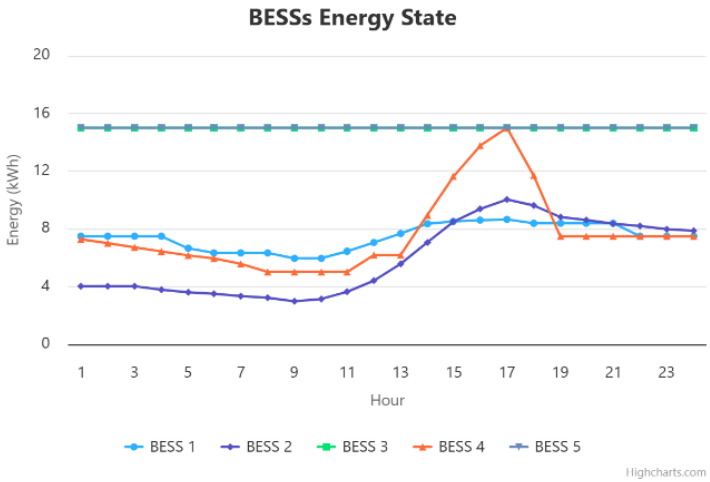

Figure 5-8: Baseline BESSs energy state

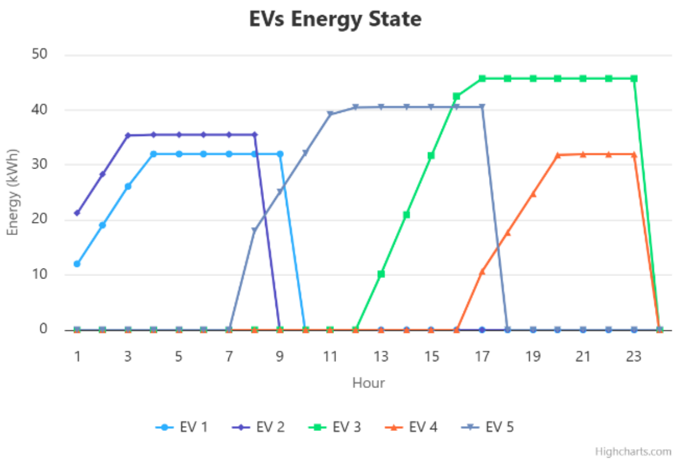

Figure 5-9: Baseline EVs energy state

5.4.2 Reduce operational cost

The reduce operational cost OF considers several factors, such as the price of each service, the charging and discharging costs of the BESSs and the EVs use based on the battery cycles, and the import and export energy prices. With this OF, grid import only occurs when there is low import price or when the internal assets are insufficient to meet demand. The largest import events occur at hours 4 and 22 (40 kW each), when the grid import price is one of the lowest from the optimisation horizon, as presented in **Figure 5-10**. Conversely, exports occur when tariffs are favourable, with the most notably export peak aligned with one of the highest export tariffs, at hour 21 with 17.42 kW (**Figure 5-11**).

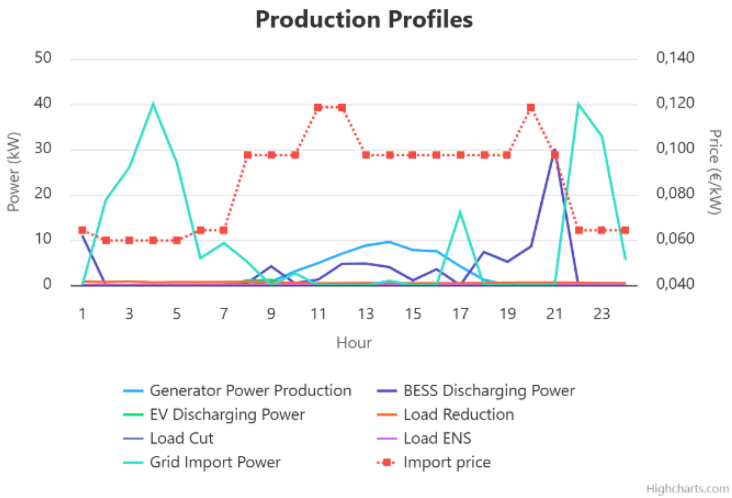

Figure 5-10: Reduce operational cost OF production profile

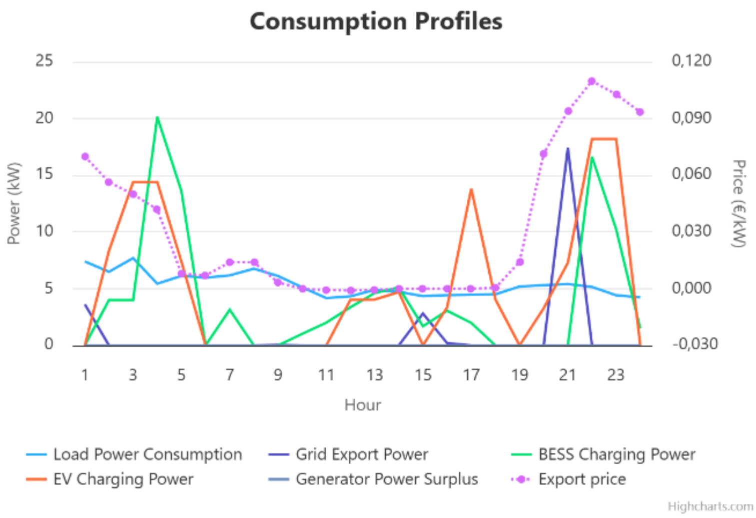

Figure 5-11: Reduce operational cost OF consumption profile

In this case demand is reduced by 9.81% (14.02kWh), yet there was no load not supplied and no load cut, as intended from the cost-minimization objective (Figure 5-12).

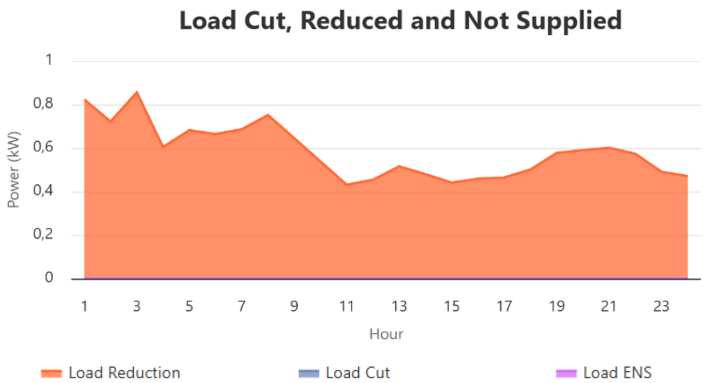

Figure 5-12: Reduce operational cost OF Load Cut, Reduced and Not Supplied

The BESSs charged overnight when import tariffs were lowest and again during peak PV generation, then discharged into the grid during high export-tariff hours (around hour 21), as presented in Figure 5-13. All EVs reached their target SoC and EV 1 employed V2X (Figure 5-14). At the end of the day, the total energy invoice of this approach is €13.26, which represents a decrease of 22.50% compared to the baseline.

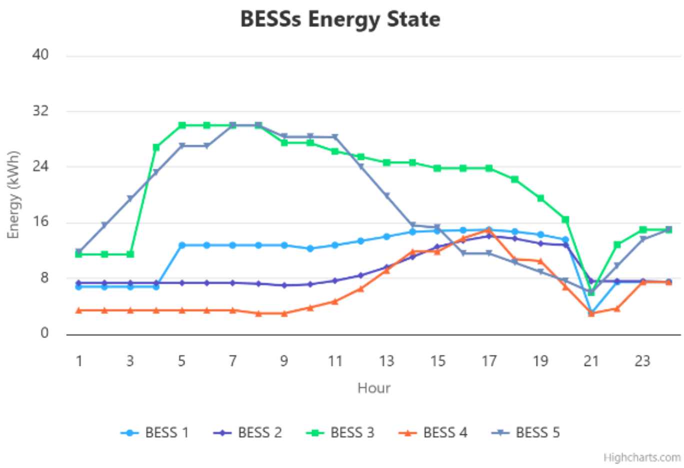

Figure 5-13: Reduce operational cost OF BESSs energy state

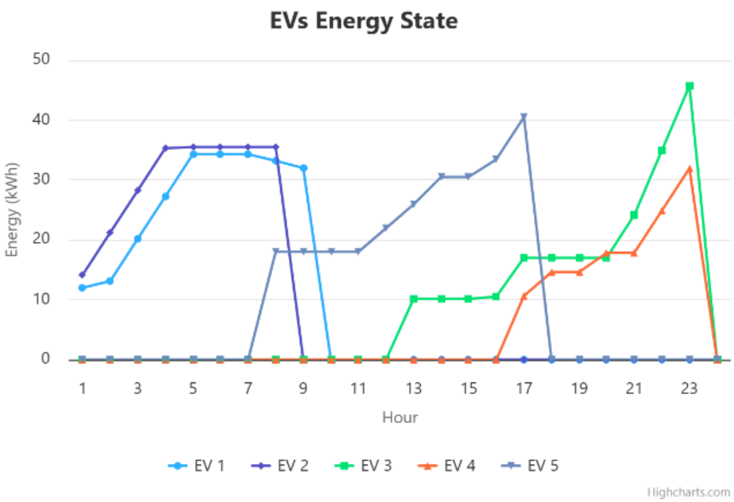

Figure 5-14: Reduce operational cost OF EVs energy state

5.4.3 Reduce energy invoice

With the reduce energy invoice OF, the objective is to minimize the total energy invoice. Results obtained show that the EC imports only occur during low-price windows or when the internal supply is insufficient to meet demand, with the largest import events occurring at hours 22 and 23 (40kW each), as presented in **Figure 5-15**. A 16.82 kW export peak occurs in hour 21, aligning with one of the highest export tariffs, evidencing invoice reduction behaviour (**Figure 5-16**).

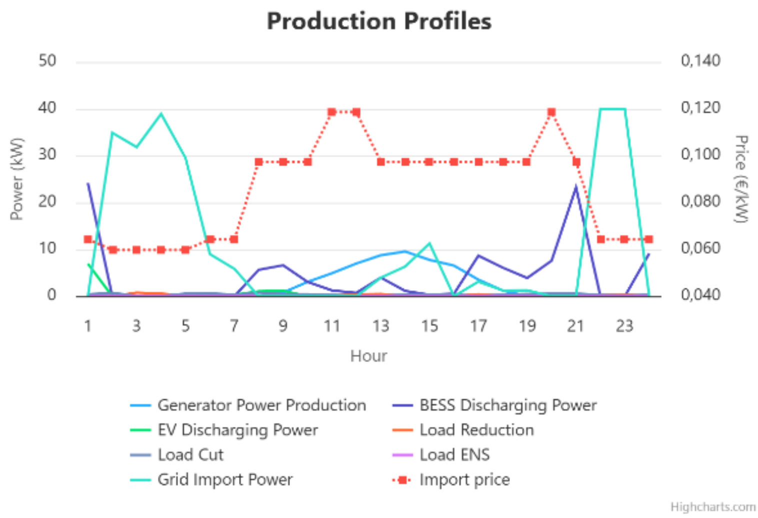

Figure 5-15: Reduce energy invoice OF production profile

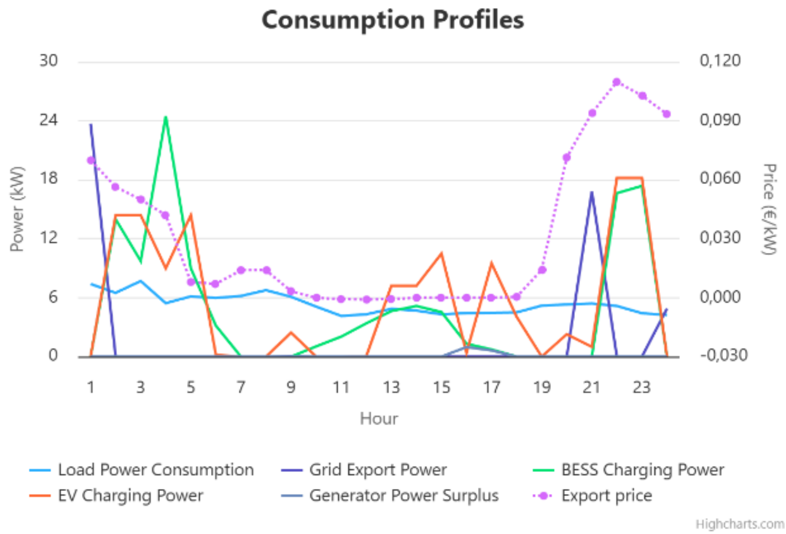

Figure 5-16: Reduce energy invoice OF consumption profile

During the optimization horizon, demand decreased by 9.96% (14.24 kWh), with 8.75 kWh achieved through load cutting and 5.49 kWh through load reduction, as presented in Figure 5-17.

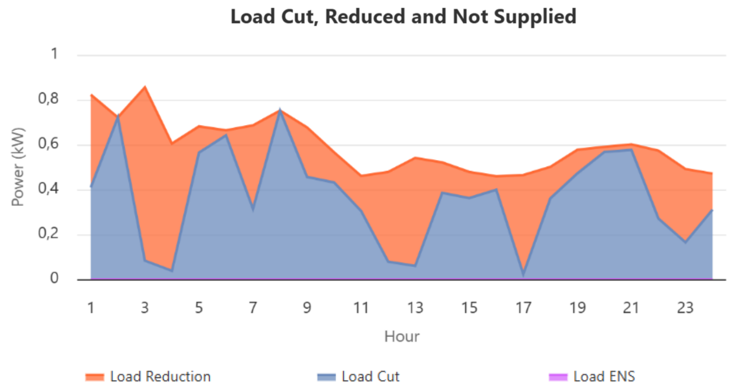

Figure 5-17: Reduce energy invoice OF Load Cut, Reduced and Not Supplied

The BESS charged overnight when import tariffs were lowest and again during peak PV generation, then discharged into the grid during high export-tariff hours, as shown on **Figure 5-18**. EV 1 employed V2X and all EVs reached their target SoC (**Figure 5-19**). Over the full day, this strategy achieves a total energy invoice of €13.13, which represents a reduction of 23.26% compared to the baseline.

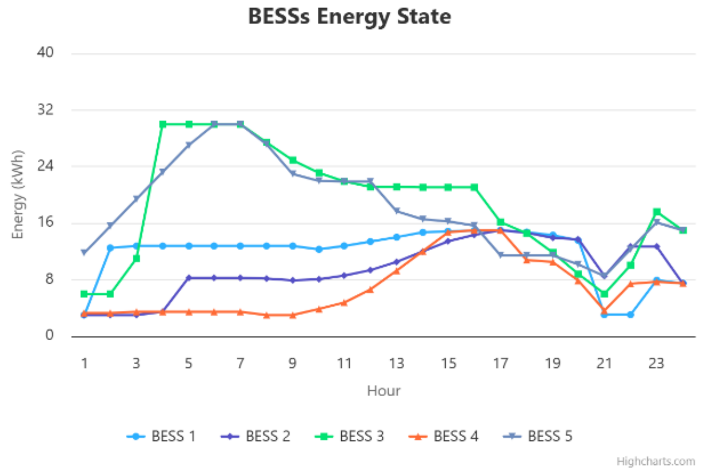

Figure 5-18: Reduce energy invoice BESSs energy state

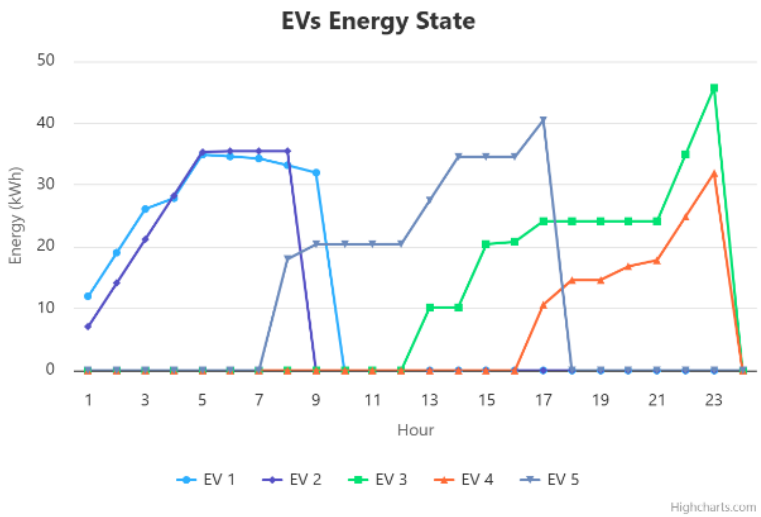

Figure 5-19: Reduce energy invoice EVs energy state

5.4.4 Reduce Grid Import

With reduce grid import OF, the objective is to minimize grid imports. Compared to the other OFs, this OF achieves the lowest level of grid import, with 166.86 kW of imported power. As shown in **Figure 5-20**, grid import only occurs during hours of high demand from EVs and loads, peaking at 20.27 kW in hour 5 when EVs 1 and 2 are charging. In parallel, the system maximizes self-consumption, resulting in no grid exports (**Figure 5-21**).

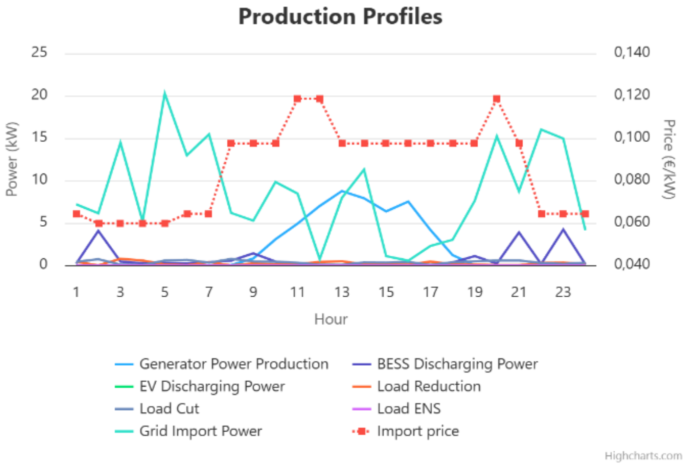

Figure 5-20: Reduce Grid Import OF production profile

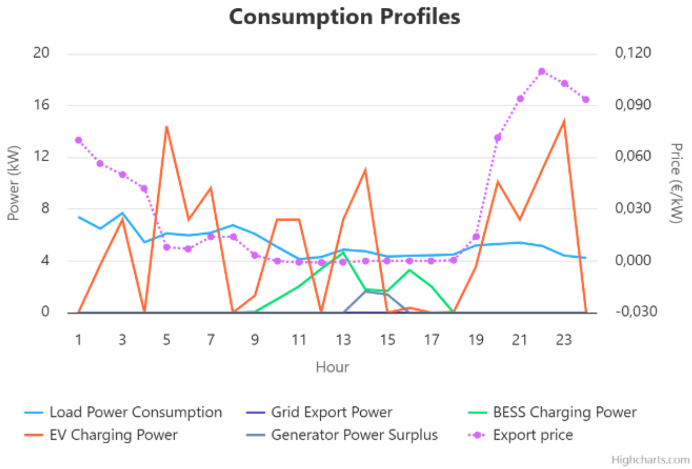

Figure 5-21: Reduce Grid Import OF consumption profile

In total, demand decreased by 9.93% (14.20 kWh). The decrease consisted in 8.70 kWh of load reduction and 5.50 kWh of load curtailment (Figure 5-22).

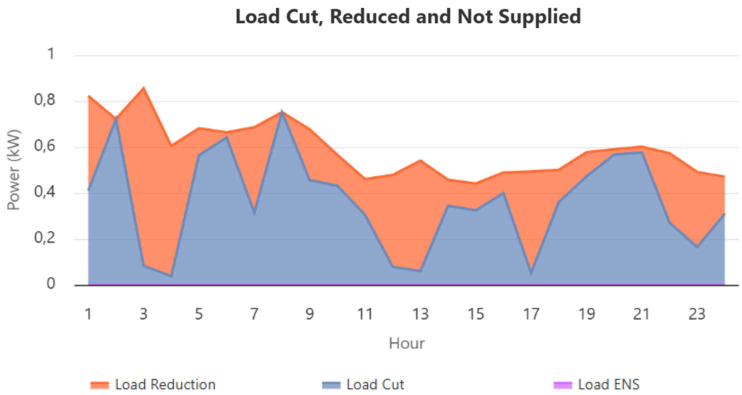

Figure 5-22: Reduce Grid Import OF Load Cut, Reduced and Not Supplied

The BESSs 1, 2 and 4 discharged overnight to power the EV's charging and then recharged during peak PV production hours (Figure 5-23). This behavior aligns with the import-minimization objective. Furthermore, all EVs reach their target SoC (Figure 5-24).

At the end of the day, the total energy invoice is €16.42, with a reduction of 4.03% compared to the baseline.

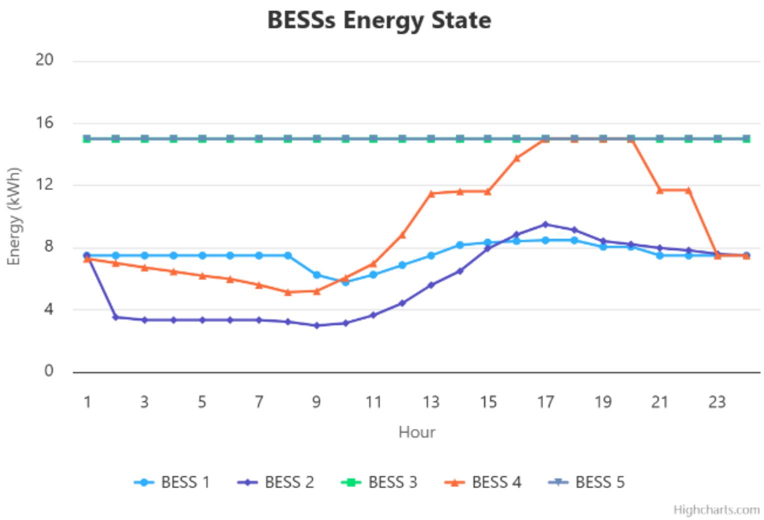

Figure 5-23: Reduce Grid Import OF BESSs energy state

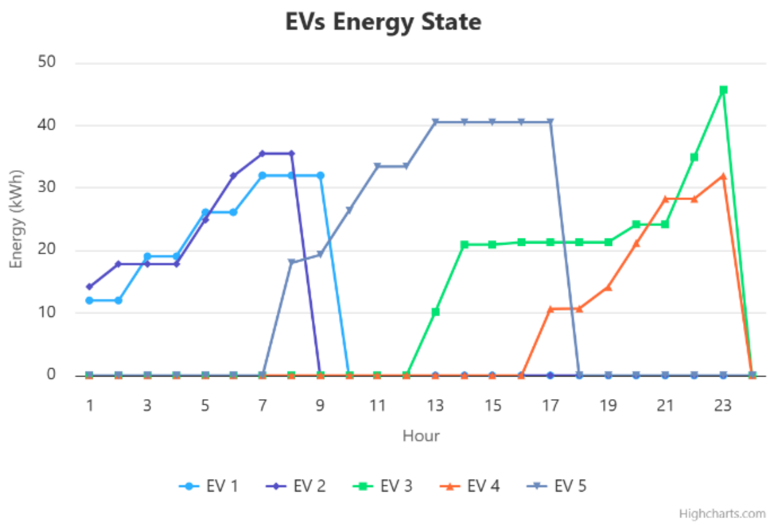

Figure 5-24: Reduce Grid Import OF EVs energy state

5.4.5 Increase Battery Longevity

With the battery longevity OF, the objective is to minimize the number of battery cycles. In this case, results obtained and presented in Figure 5-25 show that the battery remains unused throughout the entire optimisation horizon. Furthermore, all EVs reach their target SoC (Figure 5-26).

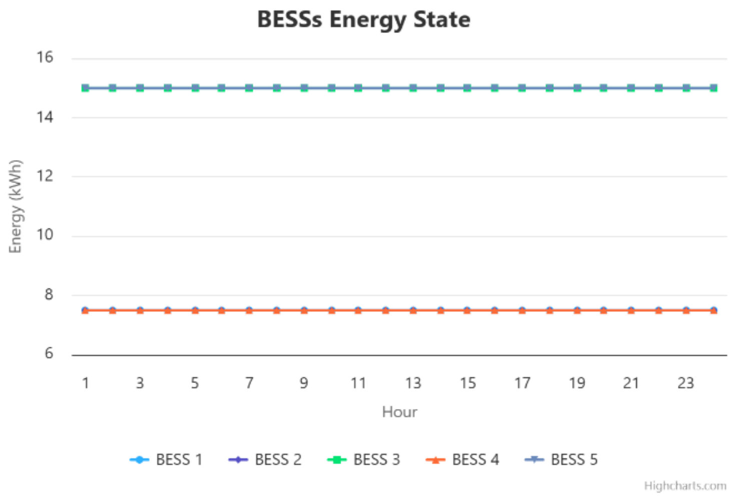

Figure 5-25: Battery longevity OF BESSs energy state

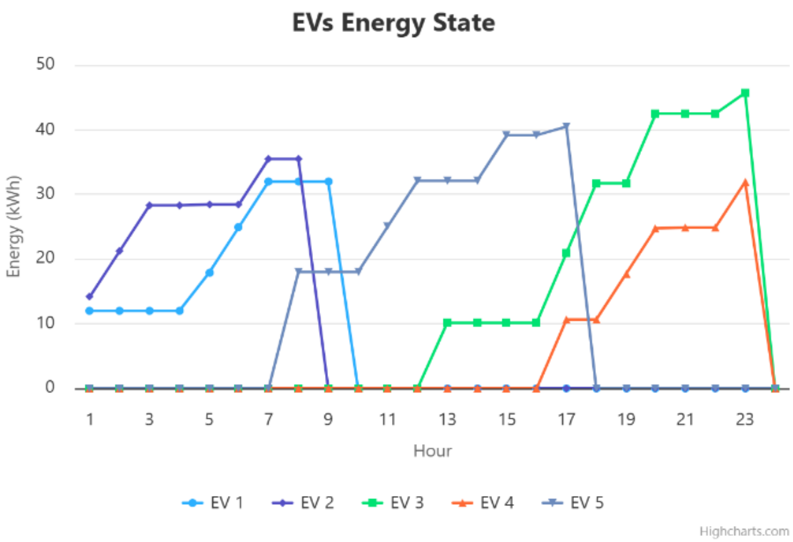

Figure 5-26: Battery longevity of EVs energy state

In **Figure *5-27*** it is shown that grid imports occur during hours without PV generation to charge the EVs and to supply the load demand. **Figure *5-28*** shows that under this OF, no grid export occurred. However, a generator power surplus was observed between hours 9 and 18, with a peak of 2.58 kW at hour 14, coinciding with high PV production. However, despite this surplus, a portion of the load still had to be curtailed during the same period.

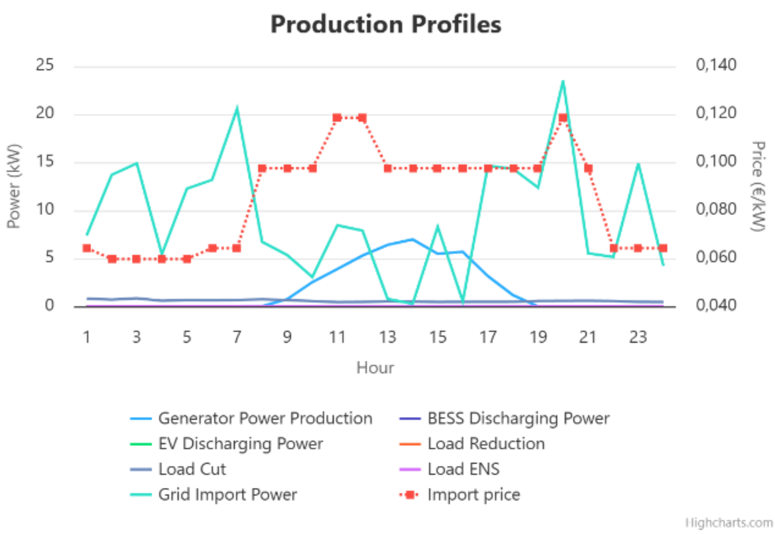

Figure 5-27: Battery Longevity OF production profile

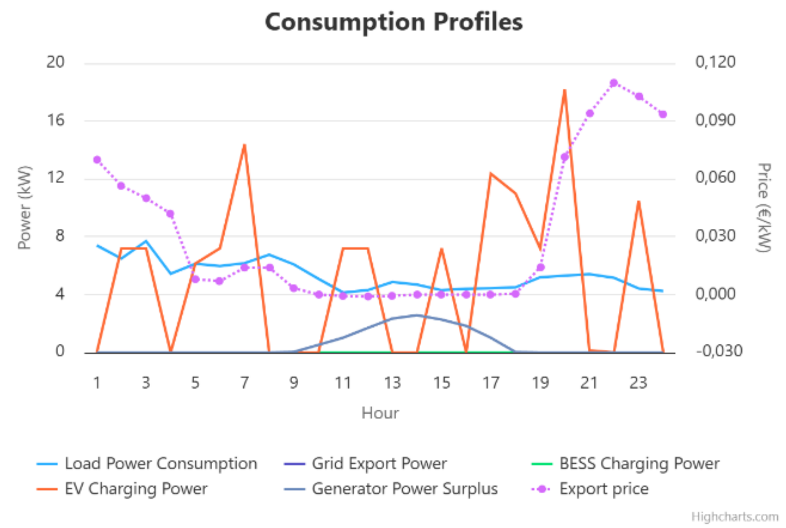

Figure 5-28: Battery Longevity OF consumption profile

In total, demand decreased by 10% (14.30kWh) and there was no load, as shown in Figure 5-29. At the end of the day, the total energy invoice of this approach is €18.72, an increase of 9.41% compared with the baseline.

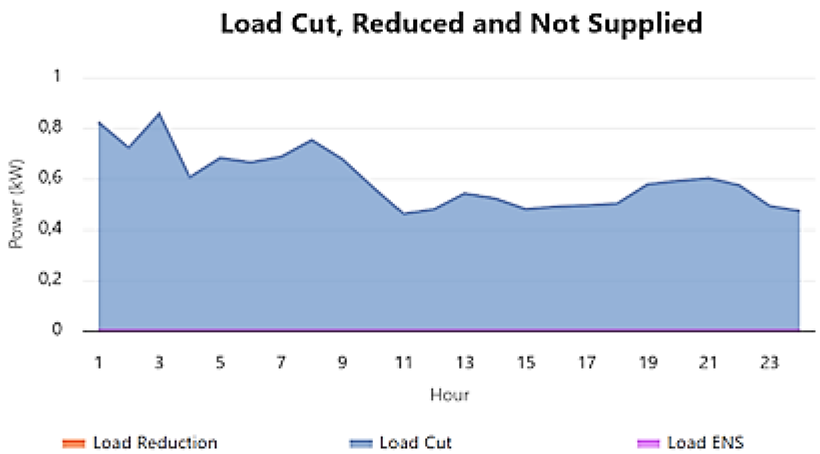

Figure 5-29: Battery Longevity OF Load Cut, Reduced and Not Supplied

5.4.6 Reduce Environmental Impact

With reduce environmental impact, the objective is to minimize CO_2 emissions by prioritizing renewable energy consumption and reducing reliance on imported electricity from non-renewable sources. As shown in **Figure 5-30**, grid imports occur mainly during hours of high demand or when the grid has a low emission factor (e.g. hour 5). Since any energy import is not entirely from renewable energy sources, this OF prioritizes self-consumption. Furthermore, no grid export was observed in this scenario, as presented in **Figure 5-31**.

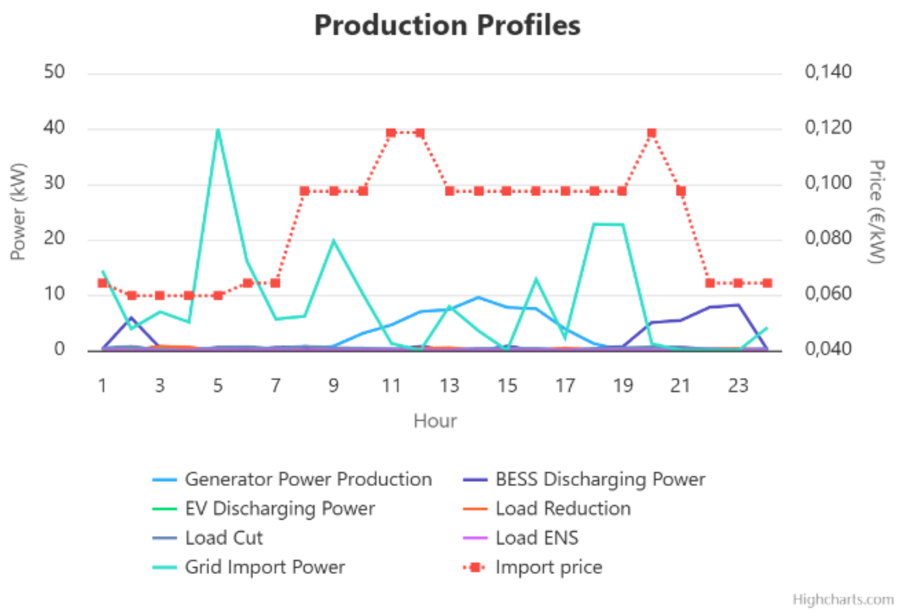

Figure 5-30: Environmental Impact OF production profile

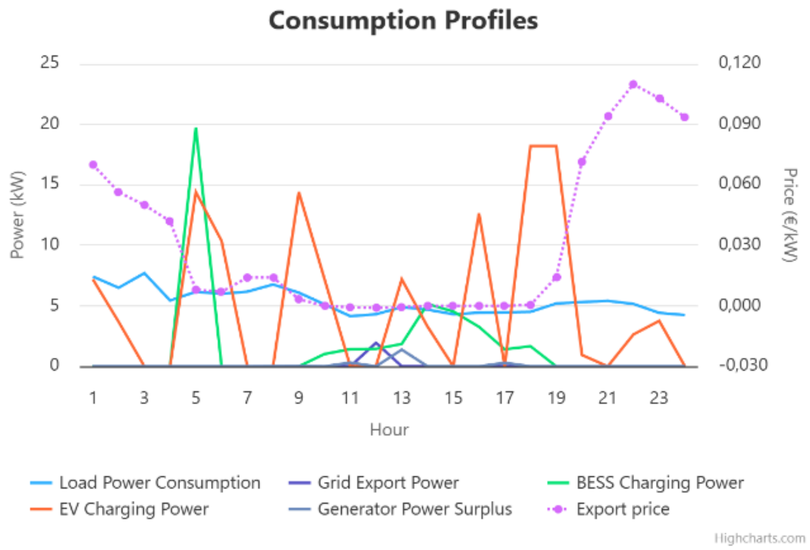

Figure 5-31: Environmental Impact OF consumption profile

From **Figure 5-32**, a grid import peak of 40 kW occurs at hour 5, coinciding with the minimum value of the emission factor. **Figure 5-33** shows that demand decreased by 9.88% (14.12kW), with a load curtailment of 8.73 kWh and a load reduction of 5.39 kWh.

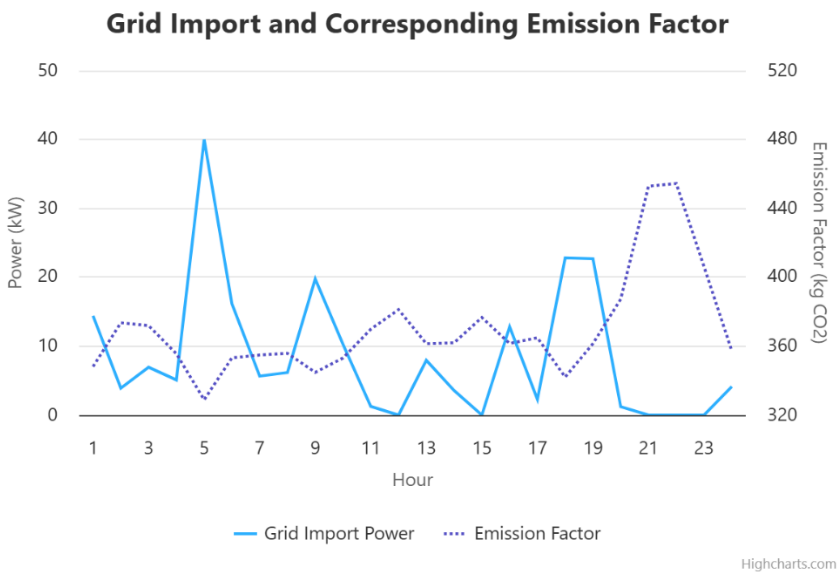

Figure 5-32: Environmental Impact OF Grid import and corresponding emission factor

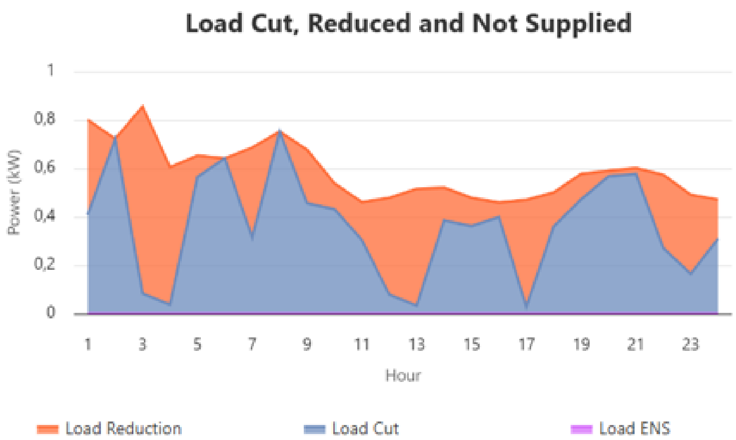

Figure 5-33: Environmental Impact OF Load Cut, Reduced and Not Supplied

During the night, the BESS discharged to support EV charging and recharged when the green factor peaked, for example units 3 and 5 between hours 4 and 5, or during peak PV generation for units 1, 2, and 4. At the end of the day, the total energy invoice of this approach is €16.80, a reduction of 1.81% compared with the baseline.

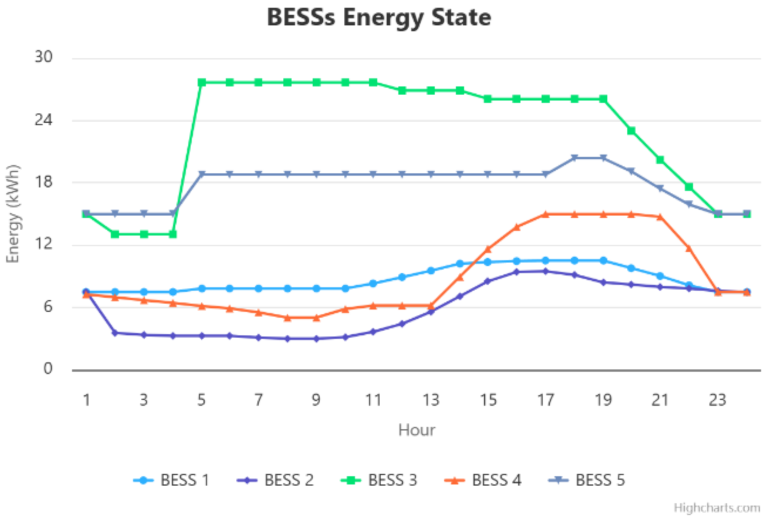

Figure 5-34: Environmental Impact OF BESSs energy state

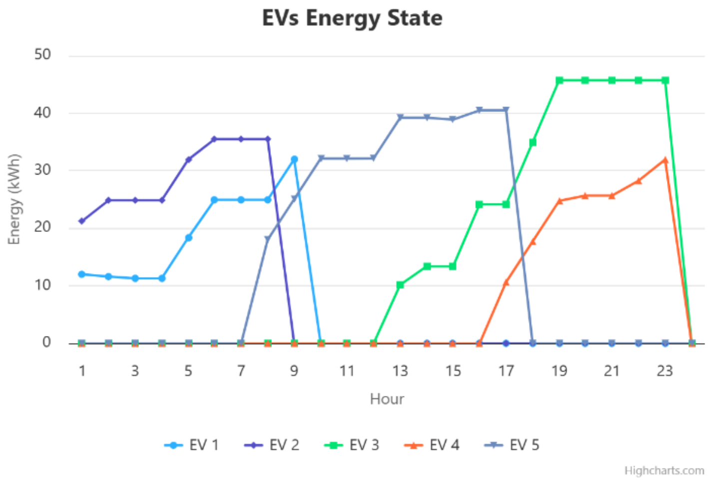

Figure 5-35: Environmental Impact OF EVs energy state

5.4.7 Increase Comfort

With the increase comfort OF, the objective is to maximize user comfort by minimizing load curtailment, load reduction and energy not supplied. Results obtained show that with this OF the resulting scheduling had no load reduction, load curtailment or load not supplied, as presented on Figure 5-36.

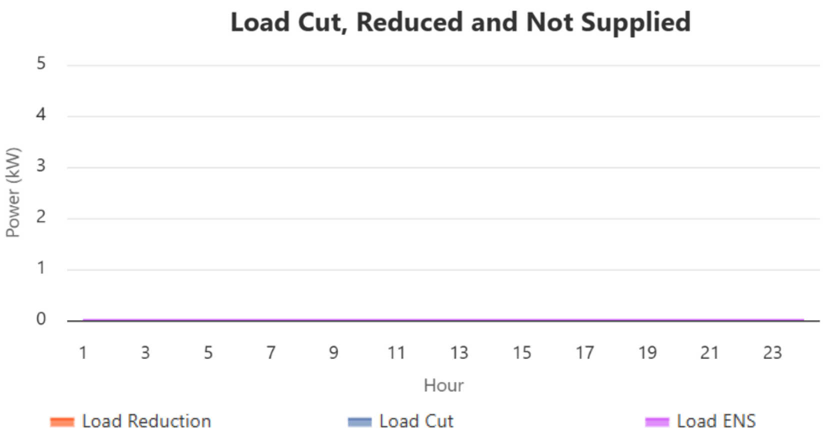

Figure 5-36: Comfort OF Load Cut, Reduced and Not Supplied

Grid imports increase to prevent demand decrease, as can be seen on Figure 5-37. Also, grid export is reduced to zero, as shown on Figure 5-38.

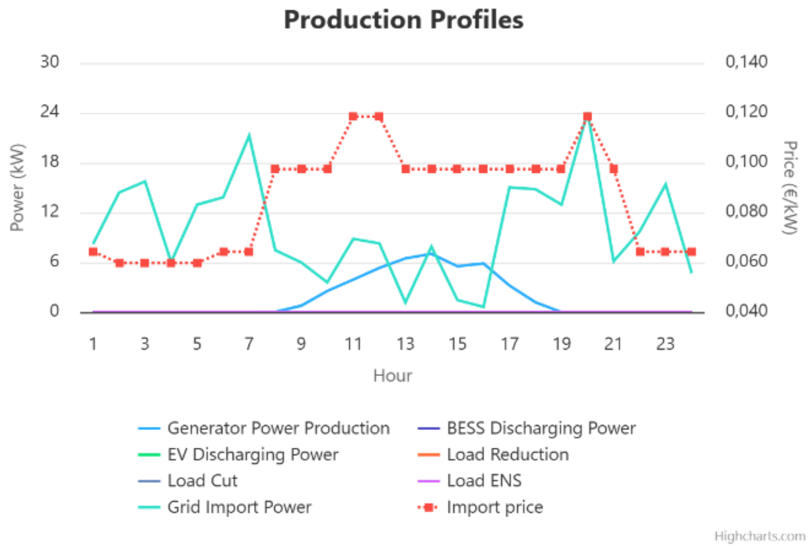

Figure 5-37: Comfort OF production profile

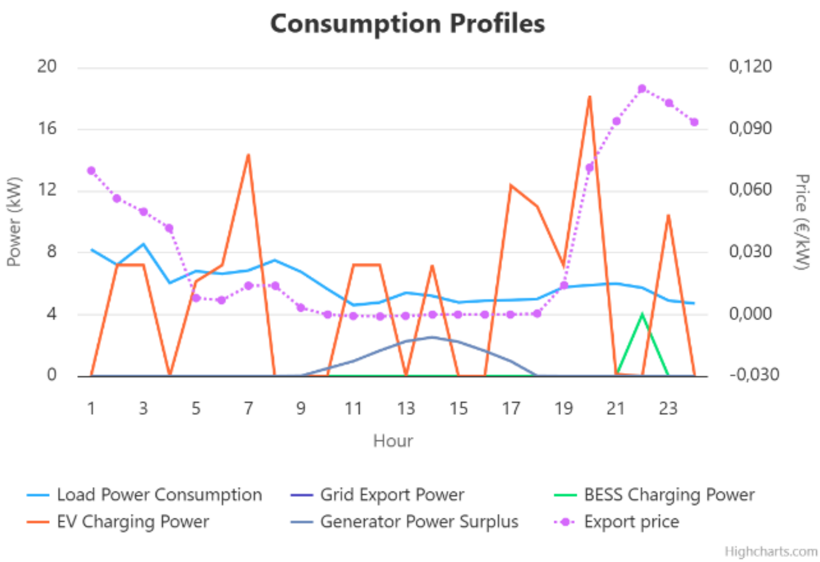

Figure 5-38: Comfort OF consumption profile

All EVs achieved their target SoC (Figure 5-40) without requiring any BESS discharge. The resulting total energy invoice for the day is €20.07, which corresponds to a 17.30% increase compared to the baseline.

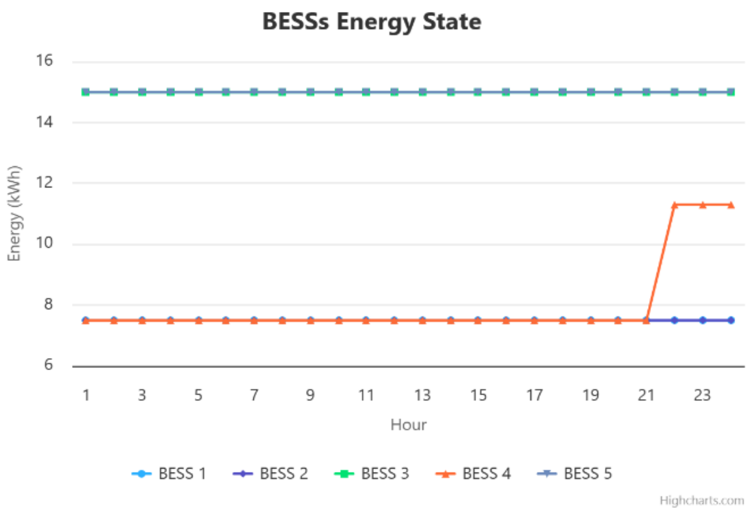

Figure 5-39: Comfort OF BESSs energy state

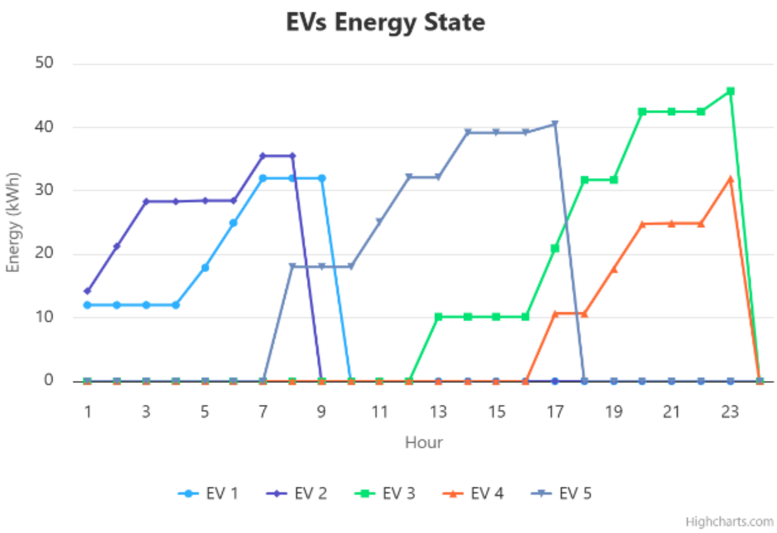

Figure 5-40: Comfort OF EVs energy state

5.5 Results analysis

Table 5-6 summarizes the KPI results for all OFs. In the table, grey highlights indicate the OF that had the best performance for each KPI, while red highlights indicate the worst performance. From this comparison, it can be observed that each OF performs best in the KPIs it specifically targets. Results obtained can be summarized as follows:

- The reduce operational cost OF had the best performance in the operational cost KPI. However, it had the worst performance on the environmental impact KPI, with an increase of 15.01% compared to the baseline OF and was the second worst in terms of battery longevity.
- The reduce energy invoice OF has the best performance in the energy invoice KPI, reaching the lowest electricity price, and having the highest level of grid export power. However, it has a worse performance regarding battery longevity and grid import and was the second worst in terms of environmental impact. This is due to the fact that the community leverages on lower import prices and high export prices to improve the performance of the scheduling w.r.t. energy invoice, while not prioritizing other factors.
- The reduce grid import OF achieves the best performance in the grid import KPI and does not perform poorly in any other KPI.
- The increase battery longevity OF achieves the best performance in the battery longevity KPI, but it has the worst performance in the total operational cost KPI, with an increase of 330.71% compared to the reduce operational cost OF and 140.62% compared to baseline OF. This outcome is explained by the fact that BESS resources are not utilized during the optimisation.
- The reduce environmental impact OF achieves the best performance in the environmental impact KPI and does not perform poorly in any other KPI.

- The increase comfort OF achieves the best performance in the comfort KPI, with no energy being cut, reduced, or left unsupplied. However, as expected, since user comfort is prioritized, the total energy invoice is the highest among all OFs, representing an increase of 17.30% compared to the baseline OF and 52.85% compared to the reduce energy invoice OF.

Table 5-6: Comparative KPI results across OFs

OF \ KPI	Energy Invoice (€)	Operational Cost (€)	Environmental Impact (kg_{CO2} / kWh)	Demand Reduction (kWh)	Battery longevity (cycles)	Grid import (kWh)	Grid Export (kWh)	V2X (Yes /No)
Baseline	17.11	34.44	0.0526	CUT = 8.80 RED = 5.13 ENS = 0 TOTAL = 13.93	2.61	204.98	0	No
Reduce Operational Cost	13.26	19.24	0.0605	CUT = 0 RED = 14.02 ENS = 0 TOTAL = 14.02	8.55	230.67	24.11	Yes
Reduce energy invoice	13.13	26.56	0.0602	CUT = 8.75 RED = 5.49 ENS = 0 TOTAL = 14.24	10.79	257.25	45.52	Yes
Reduce Grid Import	16.42	35.26	0.0543	CUT = 8.70 RED = 5.50 ENS = 0 TOTAL = 14.20	2.54	204.98	0	No
Increase Battery Longevity	18.72	82.87	0.0540	CUT = 14.30 RED = 0 ENS = 0 TOTAL = 14.30	0	223.61	0	No
Reduce Environmental Impact	16.80	31.34	0.0521	CUT = 8.73 RED = 5.39 ENS = 0 TOTAL = 14.12	3.94	207.10	1.95	Yes
Increase Comfort	20.07	80.97	0.0540	CUT = 0 RED = 0 ENS = 0 TOTAL = 0	0.27	240.83	0	No

5.6 Results – Multi Objective Function

To explore trade-offs, OFs that achieved the worst results in each KPI were combined with those that had the best performance in the same KPI. The rationale is that the advantages of one OF may counterbalance the weaknesses of another, leading to more advantageous results overall.

5.6.1 Reduce environmental impact (F6) and reduce operational cost (F2)

The reduce operational cost OF had the worst result in terms of environmental impact, with 0.0605 kg_{CO2} / kW. To address this limitation, three different combinations of the environmental impact OF with the reduce operational cost OF were evaluated.

5.6.1.1. α_6 = 0.25 and α_2 = 0.75

Combining the environmental impact OF (α_6 = 0.25) with the reduce operational cost OF (α_2 = 0.75) yields a mixed OF that emphasizes the decrease of environmental impact, while accounting for operational costs. With this OF, there is some grid export, with a peak of 4.65 kW in hour 21, which coincides with the hour of highest export price from the optimisation horizon. Battery resources are allocated efficiently, with BESSs charging during hours of lower import prices and significant PV generation, and BESSs discharging to supply EV charging and local demand.

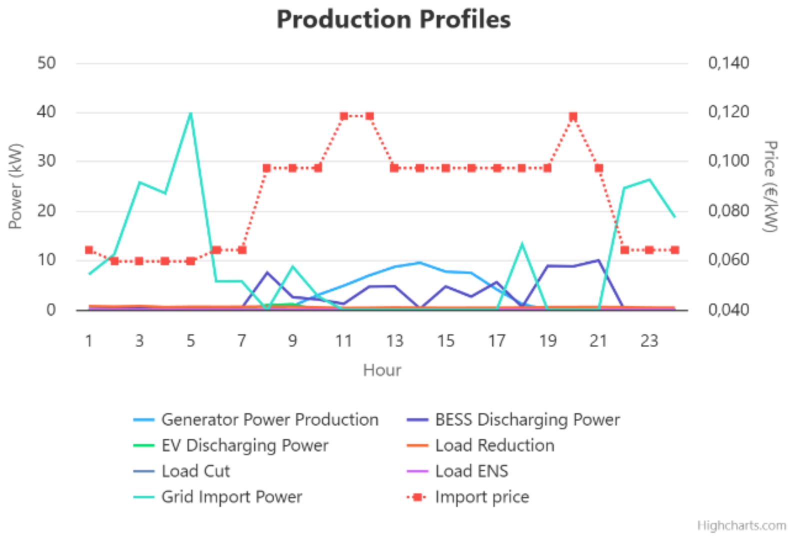

Figure 5-41: $\alpha_6 = 0.25$ and $\alpha_2 = 0.75$ OF – Production profile

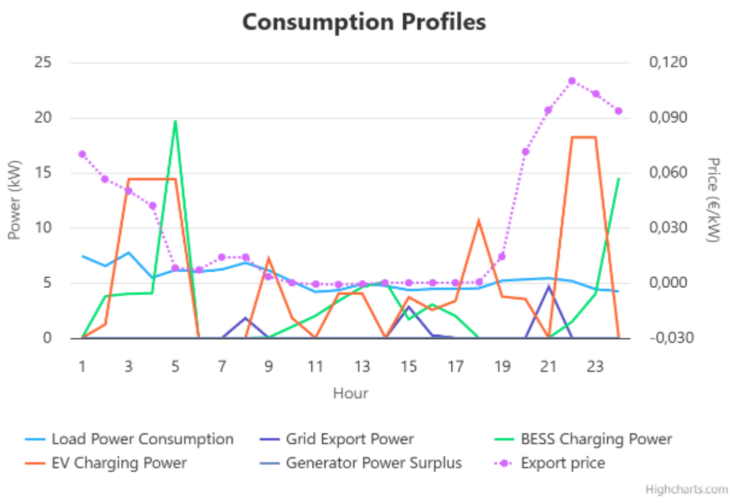

Figure 5-42: $\alpha_6 = 0.25$ and $\alpha_2 = 0.75$ OF – Consumption profile

The environmental impact factor has limited influence on the results obtained. This is evident in hour 22 of the optimisation, where grid import still occurs despite high emission factors. This energy is used by EVs 3 and 4 charging and to supply the load demand. Despite of this, by mixing both OFs, there was a decrease of 2.15% in the environmental impact KPI compared to the case where only the reduce operational cost OF was used, as presented in **Table *5-7***. Demand was reduced by 9.64% (13.77kW). Also, there was no energy cut.

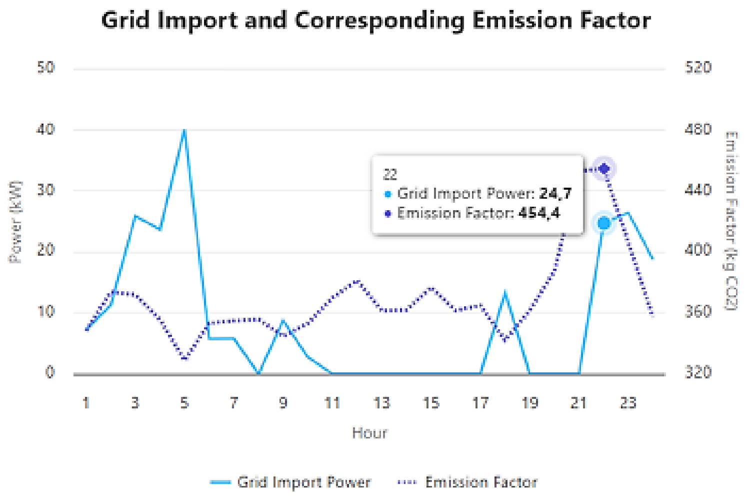

Figure 5-43: $\alpha_6 \ = \ 0.25$ and $\alpha_2 = \ 0.75$ OF – Grid Import and Associated Emissions

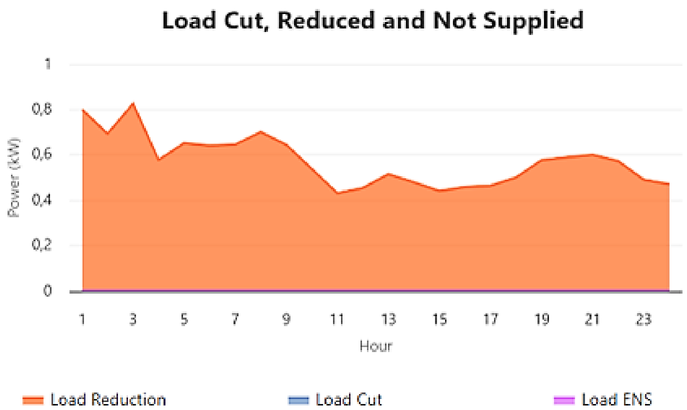

Figure 5-44: $\alpha_6 \ = \ 0.25$ and $\alpha_2 = \ 0.75$ OF – Load Cut, Reduced and Not Supplied

As presented in Figure 5-45, BESS 5 charged from hours 1 to 5, and again from hours 22 and 24, when the grid import tariff was the lowest from the optimisation horizon. BESS 3 charged around hours 4 and 5, also in periods of low grid import tariff. BESS 4 discharged through periods 1 to 7 to supply the load. Later, BESS 3 and 5 discharged, even with PV availability, to serve the load and EVs, while BESS 1, 2 and 4 charged from PV. Between periods 18 and 21, the BESS units discharged gradually to meet EV and load demand. All EVs successfully reached their target SoC (**Figure *5-46***). By the end of the day, the total energy invoice of this approach is €13.68, which represents an increase of 3.17% compared to the reduce operational cost OF.

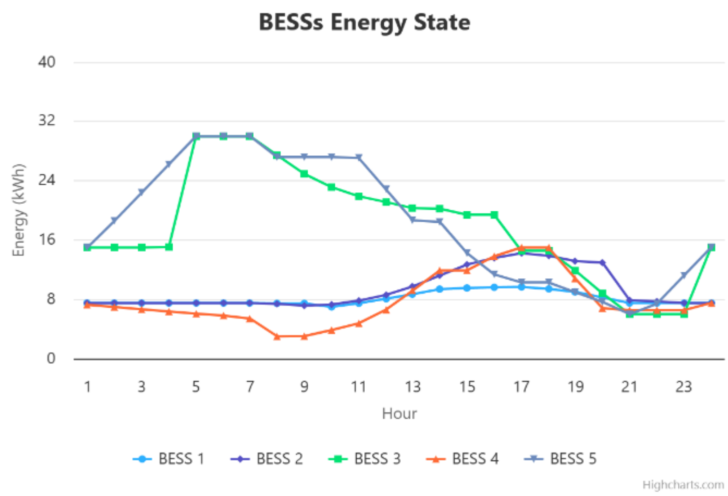

Figure 5-45: $\alpha_6 = 0.25$ and $\alpha_2 = 0.75$ OF – BESSs energy state

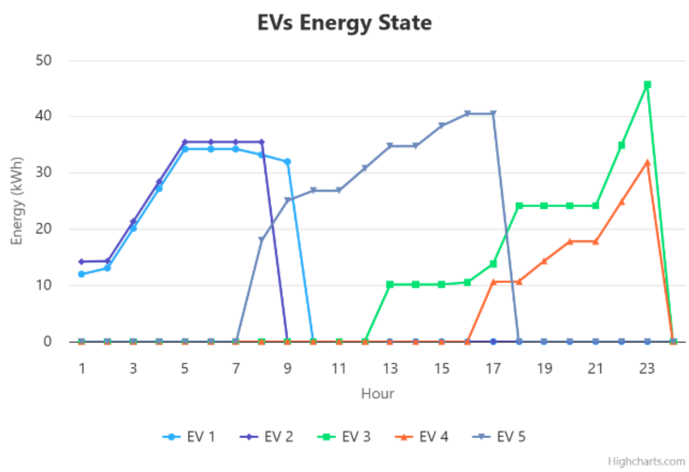

Figure 5-46: $\alpha_6 = 0.25$ and $\alpha_2 = 0.75$ OF – EVs energy state

5.6.1.2. α_6 = 0.5 and α_2 = 0.5

With equal weights on the reduce environmental impact and the reduce energy invoice objectives (α = 0.5 each), the mixed objective results in lower grid exports, as illustrated in Figure 5-47. Exports occur only in hours 8, 15, and 21. During hours 8 and 15, the tariff is close to zero, while during hour 21 the tariff is close to its maximum, showing that the export price has limited importance on the multi-objective function.

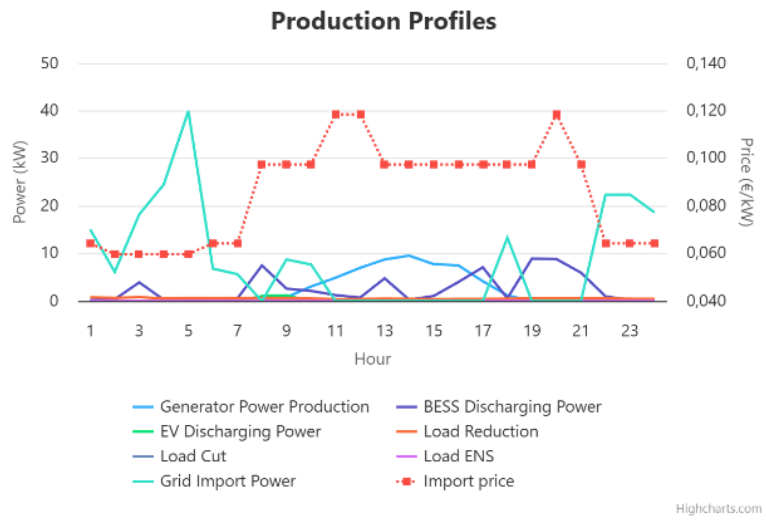

Figure 5-47: $\alpha_6 = 0.5$ and $\alpha_2 = 0.5$ OF – Production profile

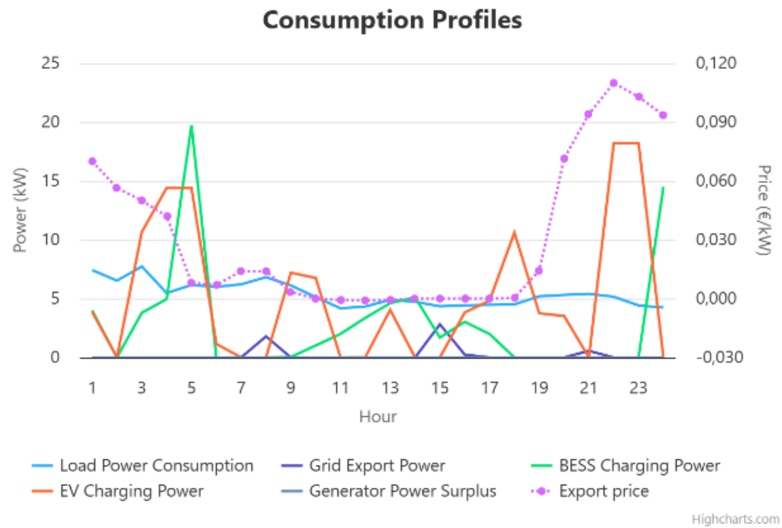

Figure 5-48: $\alpha_6 = 0.5$ and $\alpha_2 = 0.5$ OF – Consumption profile

This OF has a better performance in terms of environmental impact KPI with a decrease of 3.64% compared to the case where only the reduce operational cost OF was used, as presented in **Table *5-7***. Despite of this, the reduce environmental impact OF has limited influence on the multi-objective function, as presented in Figure 5-49. In particular, at hour 22 of the optimisation, a grid import event of 22.40 kWh occurs despite high emission factors. This behaviour is similar to that found with the export values of Figure 5-47 in the sense that, the final scheduling was not obtained while considering the reduce environmental OF alone. Although demand was reduced by 9.64% (13.77kW), no energy was left unsupplied and no load curtailment occurred, as can be seen on Figure 5-50.

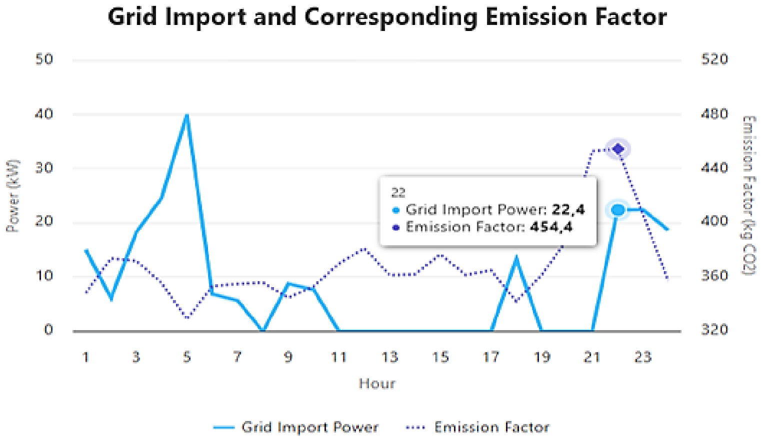

Figure 5-49: $\alpha_6 = 0.5$ and $\alpha_2 = 0.5$ OF – Grid Import and Associated Emissions

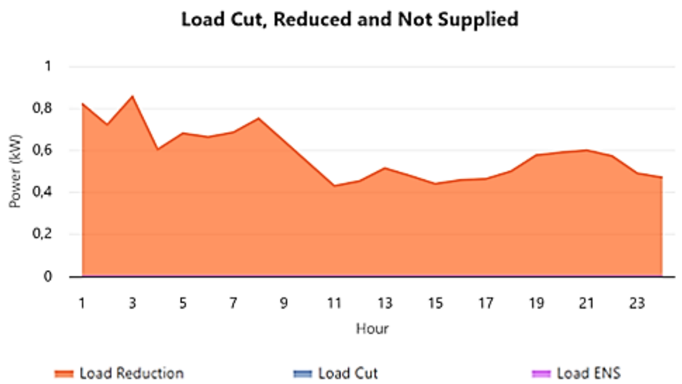

Figure 5-50: $\alpha_6 = 0.5$ and $\alpha_2 = 0.5$ OF – Load Cut, Reduced and Not Supplied

As presented in Figure 5-51, BESS 5 charged from hours 2 to 5, and again from hours 23 to 24, when the grid import tariff was the lowest from the optimisation horizon. BESS 3 charged from hours 3 to 5, also in periods of low grid import tariff. BESS 4 discharged through hours 1 to 8 to supply the load. Later, BESS 3 and 5 discharged, even with PV availability, to supply the load and EVs, while BESS 1, 2 and 4 charged from PV. Between hours 18 and 21, the BESS units discharged gradually to meet EV and load demand. All EVs reach their target SoC and EV 1 employs V2X (**Figure *5-52***). By the end of the day, the total energy invoice of this approach is €13.99, which represents an increase of 5.51% compared to the reduce operational cost OF.

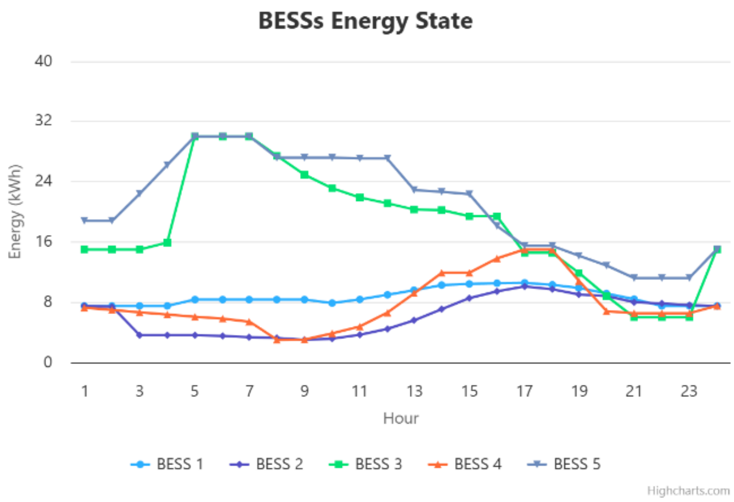

Figure 5-51: $\alpha_6 = 0.5$ and $\alpha_2 = 0.5$ OF – BESSs energy state

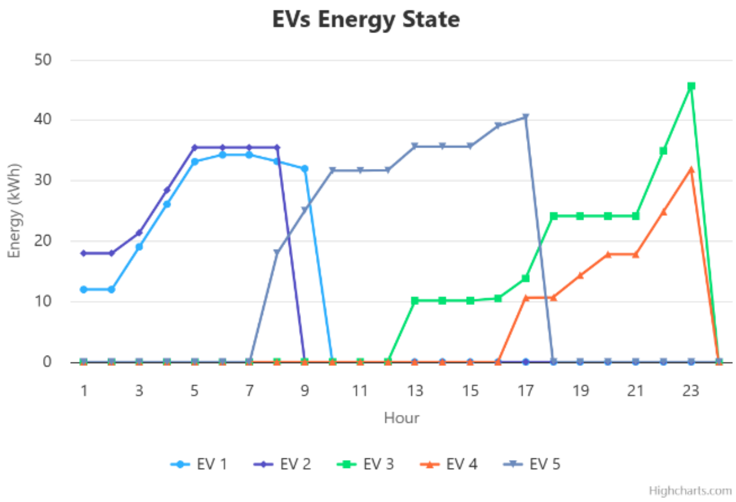

Figure 5-52: $\alpha_6 = 0.5$ and $\alpha_2 = 0.5$ OF – EVs energy state

5.6.1.3. α_6 = 0.75 and α_2 = 0.25

Combining the reduce environmental impact OF (α_6 = 0.75) with the reduce operational cost OF (α_2 = 0.25) yields a mixed objective that strongly prioritizes minimizing environmental impact. Under this objective, grid exports are largely avoided; small exports occur only in hours 15, 16, and 22. The largest export, 2.816 kW in hour 15, presented in **Figure *5-54***, coincides with a zero export tariff, so it is not economically justified, yet there is also a small energy export at period 22 with the highest export tariff, showing the low impact of the reduce operational cost.

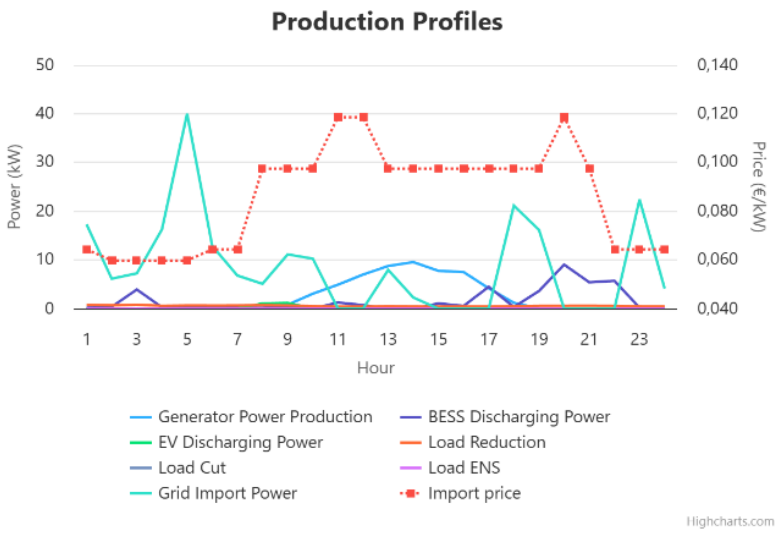

Figure 5-53: $\alpha_6 = 0.75$ and $\alpha_2 = 0.25$ OF – Production profile

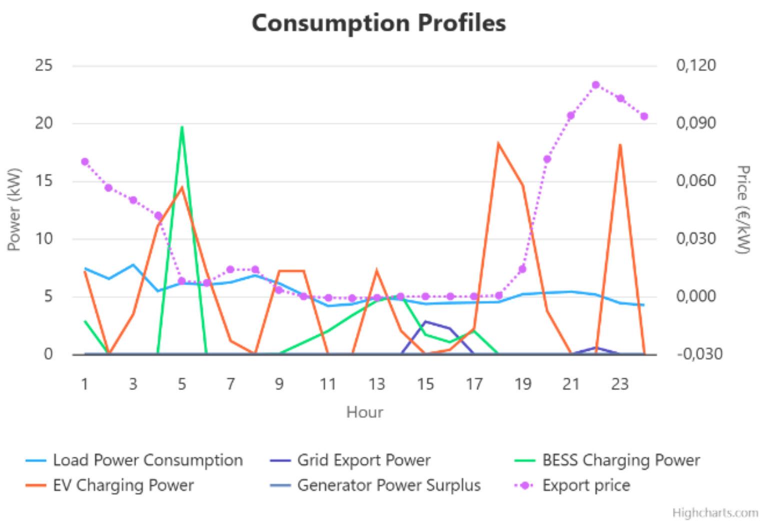

Figure 5-54: $\alpha_6 = 0.75$ and $\alpha_2 = 0.25$ OF – Consumption profile

A grid import peak of 40 kW occurs at hour 5, coinciding with the minimum value of the emission factor. With $\alpha_6 = 0.75$, there is no grid import in the hours of highest emissions, as can be seen on **Figure *5-55***. Demand was reduced by 10% (13.77 kW) as presented on Figure 5-56.

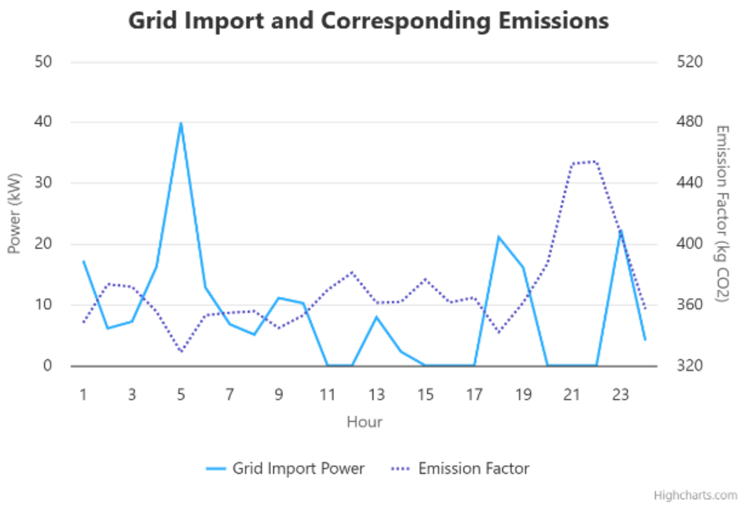

Figure 5-55: $\alpha_6 = 0.75$ and $\alpha_2 = 0.25$ OF – Grid Import and Associated Emissions

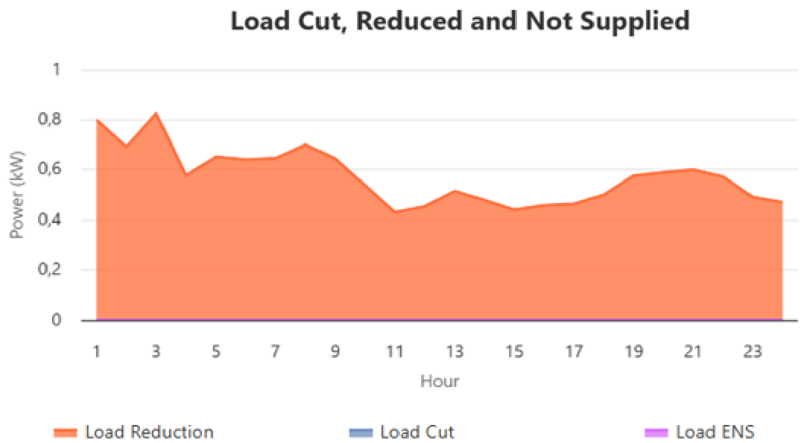

Figure 5-56: $\alpha_6 = 0.75$ and $\alpha_2 = 0.25$ OF – Load Cut, Reduced and Not Supplied

In general, as can be seen in Figure 5-57, the BESSs charged during the night, when import tariffs were lowest and again during peak PV generation, then discharged into the grid during high export-tariff hours, since the OF considers the grid export price tariff. Despite of this, BESS 3 discharged even in moments of PV generation since the EV 3 is connected during the day. All EVs reached their target SoC (**Figure *5-58***). By the end of the day, the total energy invoice of this approach is €15.41, which represents a decrease of 16.21% compared to the reduce operational cost OF.

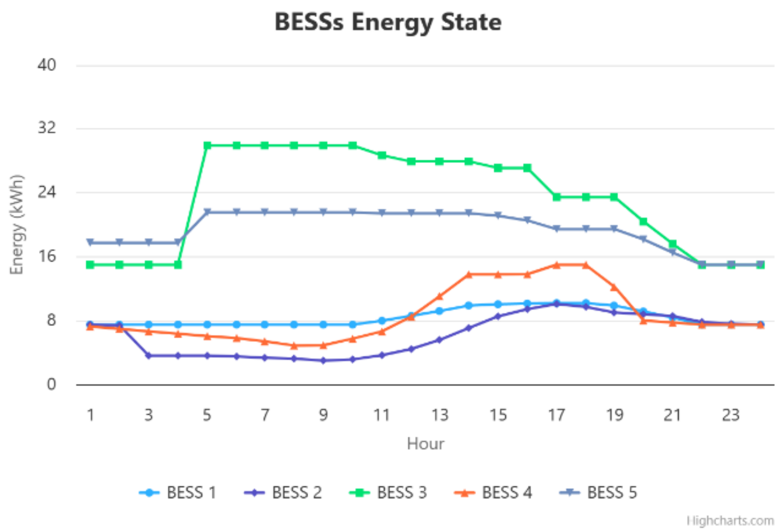

Figure 5-57: $\alpha_6 = 0.75$ and $\alpha_2 = 0.25$ OF – BESSs energy state

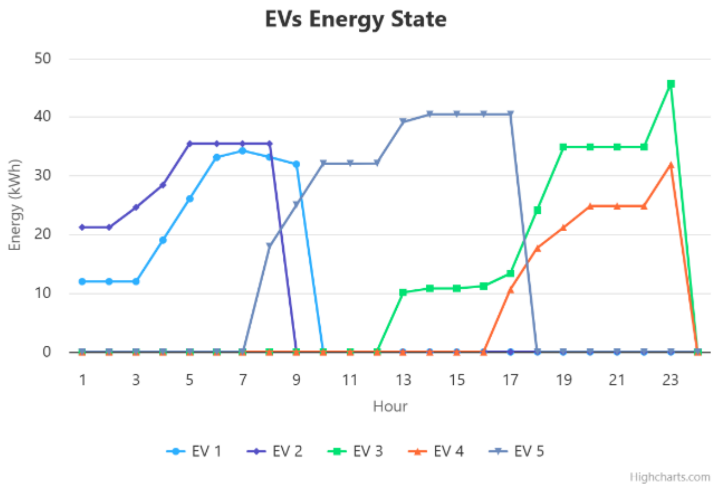

Figure 5-58: $\alpha_6 = 0.75$ and $\alpha_2 = 0.25$ OF – EVs energy state

5.6.1.4. Results Analysis for combination reduce environmental impact and reduce operational cost

Table 5-7 presents the KPIs for different combinations of reduce environmental impact OF and reduce operational cost OF. As the weight assigned to reduce environmental impact OF increases, the environmental impact, battery cycling, grid imports, and grid exports KPIs all decrease, as expected. Conversely, costs rise as the energy invoice KPI declines, which is consistent with expectations. The demand reduction KPI remains stable across all configurations.

Table 5-7: Comparative KPI results across different combinations of Reduce Environmental Impact OF and Reduce Operational Cost OF.

KPI / OF	Energy Invoice (€)	Operational cost (€)	Environmental Impact (kg_{CO2} / kWh)	Demand Reduction (kWh)	Battery longevity (cycles)	Grid import (kWh)	Grid Export (kWh)	V2X (Yes /No)
$\alpha_6 = 0$ and $\alpha_2 = 1$	13.26	19.24	0.0605	CUT = 0 RED = 14.02 ENS = 0 TOTAL = 14.02	8.55	230.67	24.11	Yes
$\alpha_6 = 0.25$ and $\alpha_2 = 0.75$	13.68	19.45	0.0592	CUT = 0 RED = 13.77 ENS = 0 TOTAL = 13.77	6.24	214.18	9.49	Yes
$\alpha_6 = 0.5$ and $\alpha_2 = 0.5$	13.99	19.72	0.0583	CUT = 0 RED = 13.77 ENS = 0 TOTAL = 13.77	6.01	209.65	5.41	Yes
$\alpha_6 = 0.75$ and $\alpha_2 = 0.25$	15.41	20.90	0.0549	CUT = 0 RED = 13.77 ENS = 0 TOTAL = 13.77	4.09	207.28	5.60	Yes
$\alpha_6 = 1$ and $\alpha_2 = 0$	16.80	31.34	0.0521	CUT = 8.73 RED = 5.39 ENS = 0 TOTAL = 14.12	3.94	207.10	1.95	Yes

Figure 5-59 presents the trade-off in environmental impact and operational cost obtained by changing the parameters of the individual OFs, F6 and F2, in the multi-objective function.

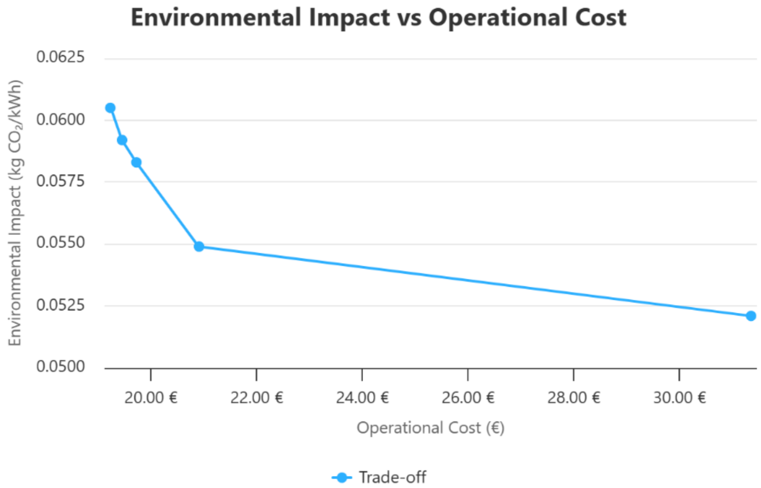

Figure 5-59: Trade-off KPIs – Environmental Impact vs Operational Cost

5.6.2 Reduce energy invoice (F3) and Increase battery longevity (F5)

The reduce energy invoice OF recorded the worst performance in terms of battery longevity, with 10.79 battery cycles used in a single day of simulation. Conversely, the battery longevity OF delivered the second-worst result in reduce energy invoice, amounting to €18.72, an increase of 42.57% in the energy invoice compared to the reduce energy invoice OF. To overcome these limitations, three different combinations of the reduce energy invoice and battery longevity OFs were evaluated to assess their combined performance.

The KPI results of the combination of the reduce energy invoice OF and the increase battery longevity OF are presented in Table 5-8. It can be observed that as the weight assigned to the reduce energy invoice OF increases, grid imports, grid exports, and the associated environmental impact also rise. This occurs because, when prioritizing costs, the OF does not account for minimizing imports and uses the grid export price to decrease the costs. Neither of the two OFs explicitly considers environmental impact; however, due to the higher level of grid imports under the reduce energy invoice OF compared to the battery longevity OF, its environmental performance is worse.

Furthermore, as the weight of the battery longevity OF decreases, a greater number of battery cycles are utilized, which contributes to lowering costs since the BESS is a valuable resource in the optimisation. Although the reduce operational cost OF accounts for both energy invoice and equipment operational costs, the share of total energy invoice remains substantially higher than the cost associated with BESS usage. Consequently, when the reduce energy invoice OF is given more weight, both the operational cost and the energy invoice decrease. Throughout the optimisation, the demand reduction KPI does not change significantly. It is also noteworthy that V2G technology is activated whenever the reduce energy invoice OF is included.

Table 5-8: Comparative KPI results across different combinations of Reduce energy invoice OF and Increase Battery Longevity OF.

KPI / OF	Energy Invoice (€)	Operational Cost (€)	Environmental Impact (kg_{CO2} / kWh)	Demand Reduction (kWh)	Battery longevity (cycles)	Grid import (kWh)	Grid Export (kWh)	V2X (Yes /No)
$\alpha_3 = 0$ and $\alpha_5 = 1$	18.72	82.87	0.0540	CUT = 14.30 RED = 0 ENS = 0 TOTAL = 14.30	0	223.61	0	No
$\alpha_3 = 0.25$ and $\alpha_5 = 0.75$	16.90	78.47	0.0554	CUT = 8.36 RED = 5.26 ENS = 0 TOTAL = 13.62	0	223.4	0	Yes
$\alpha_3 = 0.5$ and $\alpha_5 = 0.5$	15.21	37.22	0.0549	CUT = 8.79 RED = 5.71 ENS = 0 TOTAL = 14.21	2.48	211.56	4.82	Yes
$\alpha_3 = 0.75$ and $\alpha_5 = 0.25$	13.38	31.41	0.0600	CUT = 8.70 RED = 5.52 ENS = 0 TOTAL = 14.22	7.72	228.58	17.72	Yes
$\alpha_3 = 1$ and $\alpha_5 = 0$	13.13	26.56	0.0602	CUT = 8.75 RED = 5.49 ENS = 0 TOTAL = 14.24	10.79	257.25	45.52	Yes

Figure 5-60 presents the trade-off in energy invoice and battery longevity obtained by changing the parameters of the individual OFs, F3 and F5, in the multi-objective function.

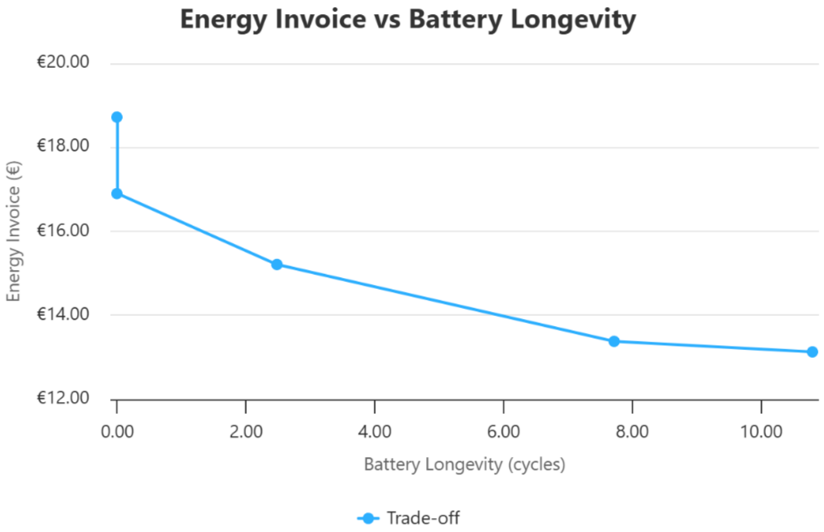

Figure 5-60: Trade-off KPIs – Energy Invoice vs Battery Longevity

5.6.3 Reduce energy invoice (F3) and Increase comfort (F7)

The increase comfort OF recorded the worst result in terms of energy invoice, leading to the highest energy invoice among all OFs, with a total of €20.07. This result is an increase of 52.86% compared to the reduce energy invoice OF, that had an energy invoice of €13.13. In contrast, the reduce energy invoice OF resulted in the second largest demand reduction, with 99.60% of the maximum permissible level of load services reduced. To mitigate these shortcomings, three alternative combinations of the reduce energy invoice and comfort OFs were tested to evaluate their joint performance.

Table 5-9 reports KPIs for different weightings of the reduce energy invoice and increase comfort objectives. As the reduce energy invoice OF weight increases, both total energy invoice and operational cost decline significantly. A notably important result was obtained with $\alpha_3 = 0.25$ and $\alpha_7 = 0.75$, because by introducing a small weight on the reduce energy invoice OF, energy invoice fell by 29.05% and equipment-related operating costs by 67.33%, without the need for load curtailment and load reduction services. This result was primarily obtained due to the increase in battery cycles from, which increased from 0.27 to 10.76 cycles over the optimisation horizon. Nevertheless, in configurations that have less weight in the increase comfort OF and more on the reduce energy invoice OF, the demand reduction increases substantially. Conversely, environmental impact KPI and both grid imports and exports KPIs increase as the comfort weight is reduced.

Table 5-9: Comparative KPI results across different combinations of reduce energy invoice OF and increase comfort OF.

OF \ KPI	Energy Invoice (€)	Operational Cost (€)	Environmental Impact (kg_{CO2} / kWh)	Demand Reduction (kWh)	Battery longevity (cycles)	Grid import (kWh)	Grid Export (kWh)	V2X (Yes /No)
α_3 = 0 and α_7 = 1	20.07	80.97	0.0540	CUT = 0 RED = 0 ENS = 0 TOTAL = 0	0.27	240.83	0	No
α_3 = 0.25 and α_7 = 0.75	14.24	26.45	0.0605	CUT = 0 RED = 0 ENS = 0 TOTAL = 0	10.76	270.88	45.16	Yes
α_3 = 0.5 and α_7 = 0.5	13.62	23.53	0.0610	CUT = 4.34 RED = 1.99 ENS = 0 TOTAL = 6.33	10.76	264.89	46.69	Yes
α_3 = 0.75 and α_7 = 0.25	13.13	24.40	0.0609	CUT = 8.56 RED = 5.46 ENS = 0 TOTAL = 14.02	10.79	257.25	46.27	Yes
α_3 = 1 and α_7 = 0	13.13	26.56	0.0602	CUT = 8.75 RED = 5.49 ENS = 0 TOTAL = 14.24	10.79	257.25	45.52	Yes

Figure 5-61 presents the trade-off in energy invoice and comfort obtained by changing the parameters of the individual OFs, F3 and F7, in the multi-objective function.

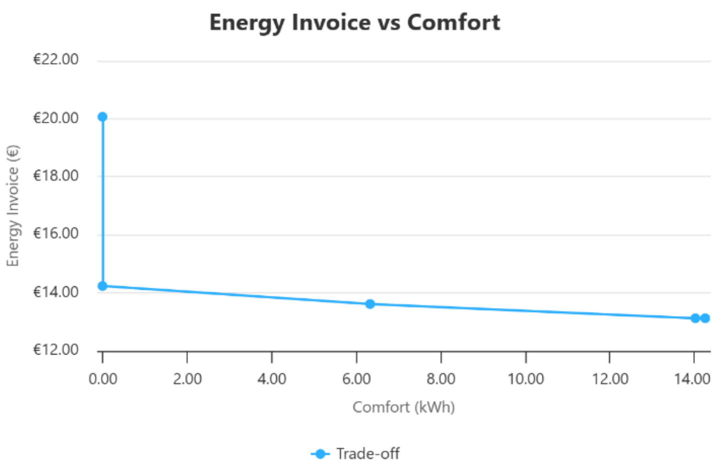

Figure 5-61: Trade-off KPIs – Energy Invoice vs Comfort

5.6.4 Reduce grid import (F4) and Reduce energy invoice (F3)

The reduce energy invoice OF had the lowest energy invoice result but the worst grid-import performance, requiring an import of 257.25 kWh from the grid. By contrast, the reduce grid import OF does not have the worst result in the energy invoice KPI, but still there is an increase in the energy invoice of 25.05% relative to the reduce energy invoice OF alone. To mitigate this trade-off, three weighted combinations of the two objectives were evaluated to assess their joint performance.

Table 5-10 shows the results across different combinations of reduce grid import OF and reduce energy invoice OF. As expected, as the reduce grid import weight increases, grid imports decline. Also, since the environmental impact KPI is linked to energy drawn from the grid, it decreases as the reduce grid import OF weight increases. Grid exports KPI also decrease, because available resources are redirected to supply the EC's internal demand. Consequently, battery cycling falls, and fewer EVs employ V2X. Conversely, the energy invoice increases when the reduce energy invoice OF weight is reduced.

Table 5-10: Comparative KPI results across different combinations of reduce grid import OF and reduce energy invoice OF.

KPI / OF	Energy Invoice (€)	Operational Cost (€)	Environmental Impact (kg_{CO2} / kWh)	Demand Reduction (kWh)	Battery longevity (cycles)	Grid import (kWh)	Grid Export (kWh)	V2X (Yes /No)
α_4 = 0 and α_3 = 1	13.13	26.56	0.0602	CUT = 8.75 RED = 5.49 ENS = 0 TOTAL = 14.24	10.79	257.25	45.52	Yes
α_4 = 0.25 and α_3 = 0.75	13.59	26.54	0.0595	CUT = 8.73 RED = 5.46 ENS = 0 TOTAL = 14.20	6.80	215.12	6.89	Yes
α_4 = 0.5 and α_3 = 0.5	13.73	24.83	0.0596	CUT = 8.70 RED = 5.56 ENS = 0 TOTAL = 14.26	6.77	210.85	3.56	Yes
α_4 = 0.75 and α_3 = 0.25	14.29	26.86	0.0580	CUT = 8.68 RED = 5.46 ENS = 0 TOTAL = 14.14	4.91	208.15	2.67	Yes
α_4 = 1 and α_3 = 0	16.42	35.26	0.0543	CUT = 8.70 RED = 5.50 ENS = 0 TOTAL = 14.20	2.54	204.98	0	No

Figure 5-62 presents the trade-off in energy invoice and grid import obtained by changing the parameters of the individual OFs, F4 and F3, in the multi-objective function.

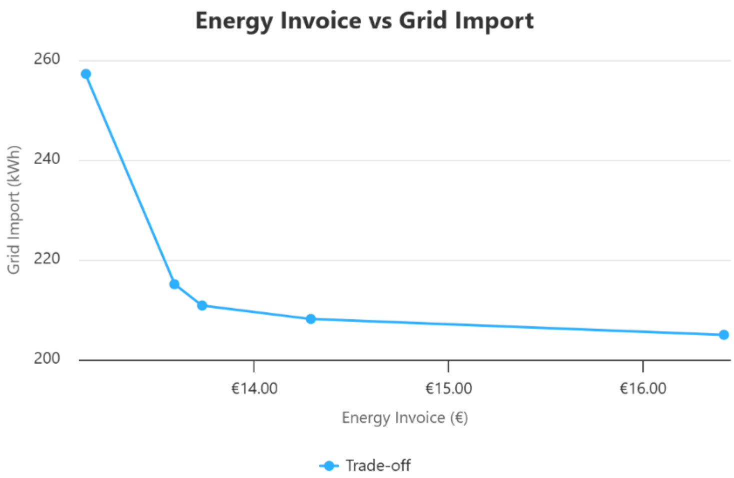

Figure 5-62: Trade-off KPIs – Grid Import vs Electricity Cost

5.7 Conclusions

In this chapter, decision support methods are proposed to assist the EC manager in the scheduling phase of internal flexible devices by incorporating individual EC member consumption preferences into a multi-objective optimization. To do that, EC member consumption preferences were first identified and modelled as individual objective functions from which EC members could choose one or multiple. Examples of such preferences include reduce energy invoice, reduce grid import, increase battery longevity, reduce environmental impact, increase comfort, and reduce operational costs.

Two main tests were considered. The first test considered that all EC members had only one OF and that the OF was the same for all EC members. Results obtained for each OF show that the resulting scheduling has the best performance for the KPIs it specifically targets. In this test, the KPIs on which the scheduling had the worst performance were also identified. The second test aimed at improving KPI performance by mixing OFs with the best and worst performances for a specific KPI. Results obtained show that it was possible, as expected, to improve KPI performance with a multi-objective optimization.

6 DER Valorisation methods

6.1 Objective and Scope

The increasing penetration of distributed energy resources (DERs) such as solar PV, batteries, electric vehicles, and flexible loads is transforming the way energy communities (ECs) operate. While these resources bring significant potential for decarbonisation and consumer empowerment, their integration also creates new challenges related to grid stability, variability of renewable generation, and efficient coordination of flexibility. Harnessing this potential requires advanced optimisation and decision-support tools that can manage the operation of heterogeneous DERs in a coordinated and economically viable way. The DER Valorisation Tool aims at providing a decision-support framework that enables energy communities (ECs) and active consumers to optimally utilize the flexibility of distributed energy resources (DERs). It leverages the spatio-temporal distribution of DERs and demand-response resources to maximise their penetration in the electricity grid while unlocking both operational and economic value. The tool provides decision-support capabilities for the optimal utilization of DER flexibility, covering:

- Renewable generation (e.g., rooftop solar PV).
- Energy storage systems (stationary batteries).
- Electric vehicles (EVs) with dynamic availability constraints.
- Flexible loads including HVAC systems and demand-response appliances.

The scope of the tool is to perform a multi-period optimisation of DER operation over a defined time horizon (e.g., 24 hours with hourly resolution). By considering forecasts of demand and renewable generation, the tool determines the optimal scheduling of DERs while ensuring:

- Feasible flexibility limits for upward and downward power injections into the grid.
- Satisfaction of all technical and operational constraints of DERs (state of charge dynamics, availability schedules, comfort bounds, generation limits, etc.).
- Maximisation of DER valorisation, i.e., improving self-consumption of local renewable generation, minimising curtailment, and creating additional economic value through participation in flexibility markets and P2P trading.

Concretely, the tool seeks to balance three major objectives: (i) Economic efficiency by reducing energy costs, cycling costs, and valorising flexibility in market mechanisms; (ii)Technical feasibility by respecting all system and device-level constraints, ensuring grid balance, and supporting higher DER penetration; (iii) User comfort and sustainability by maintaining thermal comfort in buildings, ensuring EVs meet mobility requirements, and maximising renewable utilisation.

6.2 Methodology

The optimisation problem is formulated as a multi-period optimisation (MPO) that determines the desired power setpoints of DERs and their maximum feasible flexibility ranges. The model accounts for the physical and operational characteristics of all DERs, including time-coupled dynamics for batteries (BS), EVs, and HVAC systems. The optimisation is carried out over a discrete planning horizon, defined by the simulation timestep Δt, the total simulation time T= [0, T], and the horizon H = [t_H, t_H+(H-1) Δt] ⊂ T, with H representing the number of timesteps. The decision-making interval is assumed to be equal to Δt. The net active power is expressed as:

$$P = P_{generation} - P_{load} \tag{6-1}$$

where $P > 0$ indicates net generation (e.g., PV output or discharge of storage) and P < 0 indicates net consumption (e.g., load demand, charging of storage). The outputs of the MPO include:

- Optimal setpoints for PV, battery charging/discharging, EV operation, HVAC demand, and other flexible loads.
- Flexibility ranges (P^{UP}, P^{down}) at each timestep, quantifying the extent of feasible adjustments around the setpoint.
- Feasible operating schedules that balance renewable utilisation, demand satisfaction, and comfort constraints while minimising curtailment and dependency on the external grid.

The approach ensures that DERs are orchestrated in a coordinated manner, enabling both grid support (through flexibility provision) and economic valorisation (through optimal dispatch and market participation).

6.2.1 DER Modelling

The multi-period optimisation requires detailed mathematical representations of the DERs participating in the flexibility framework. Each DER is modelled through its operational constraints, physical limits, and inter-temporal dynamics. The following subsections describe the models adopted for batteries, EVs, HVAC systems, PV generation, and flexible demand.

6.2.1.1. Battery Storage (BS)

Battery systems are characterised by their SoC dynamics and charging/discharging power limits. The SoC evolves across timesteps depending on charging/discharging decisions, round-trip efficiency, and degradation parameters. The SoC dynamics of the battery is given by:

$$SoC_i^{BS}(t+1) = \left(1 - \delta_i^{BS}\right) SoC_i^{BS}(t) - \frac{P_i^{BS}(t) \Delta t \eta_i^{BS}}{\bar{E}_i^{BS}} \tag{6-2}$$

$$\underline{P_i^{BS}} < P_i^{BS}(t) < \overline{P_i^{BS}} \tag{6-3}$$

$$\underline{SoC_i^{BS}} < SoC_i^{BS}(t) < \overline{SoC_i^{BS}} \qquad (6\text{-}4)$$

where δ_i^{BS}, η_i^{BS}, $\bar{E}_i^{BS}$ are the BS self-discharge rate, roundtrip efficiency, and maximum capacity, respectively. During BS operation, the prosumer aims to minimize the cycling cost to avoid excessive charge and discharge cycles, which can degrade the battery's lifetime.

$$f_i^{BS}(P_i^{BS}) = \alpha_{cyc} \sum_{t_H}^{t_H+(H-1)\Delta t} (P_i^{BS}(t+1) - P_i^{BS}(t))^2 \qquad (6\text{-}5)$$

where the summation extends over all time steps within the planning horizon starting at time t_H and the coefficient $\alpha_{cyc}$is a weighting factor that penalizes abrupt power variations, thereby discouraging frequent charging and discharging cycles to preserve battery health.

6.2.1.2. Electric Vehicles (EVs)

EVs are modelled similarly to stationary batteries, with additional availability constraints to reflect connection times. We assume EVs are not present at the building or home during the period $[t_1, t_2]$, e.g. when the owner is at work.

$$P_i^{EV}(t) = 0, \forall\, t \in [t_1, t_2] \qquad (6\text{-}6)$$

We include a similar cycling cost term for the EVs to limit excessive charging and discharging, thereby mitigating battery degradation and extending its lifetime. In addition, a tracking objective is introduced to reflect user mobility requirements: the EV must reach a predefined state of charge (SoC) by a specified target time t^* for example, ensuring the battery is 80% charged before commuting. Together, these terms define the EV-specific contribution to the overall optimisation objective. Thus, the EV objective function is given by:

$$f_i^{EV} = \alpha_{cyc} \sum_{t_H}^{t_H+(H-1)\Delta t} (P_i^{EV}(t+1) - P_i^{EV}(t))^2 + \xi_{EV}(SoC_i^{EV}(t^*) - SoC_i^{EV^*})^2 \qquad (6\text{-}7)$$

Where $P_i^{EV}(t)$ denotes the EV charging/discharging power at time t, $SoC_i^{EV}(t)$ is the EV state of charge at time t, $SoC_i^{EV^*}$ is the target state of charge, α_{cyc} represents the weight for EV cycling cost, ξ_{EV} the weight for meeting the mobility-related SoC target and t^* desired time by which the target SoC must be achieved.

6.2.1.3. HVAC Systems

HVAC units are modelled using a simplified thermal dynamics equation that relates indoor temperature to outdoor temperature, HVAC power consumption, and the thermal inertia of the building. This representation captures how power drawn by the unit directly impacts indoor comfort levels. In this framework, the HVAC system is assumed to operate as a heat pump (HP), capable of functioning either in heating or cooling mode depending on the ambient temperature.

- When the outdoor temperature $T_i^{out}$is higher than the indoor temperature T_i^{in}, the heat pump operates in cooling mode (acting as an air conditioner).
- Conversely, when T_i^{out} is lower, the unit provides heating to maintain comfort.

The indoor temperature dynamics are expressed as:

$$T_i^{in}(t+1) = \theta_i T_i^{in}(t) + (1-\theta_i)(T_i^{out}(t) \pm \rho_i P_i^{HP}(t)) \quad (6\text{-}8)$$

Where T_i^{in} denotes the indoor air temperature at time t, T_i^{out} represents the outdoor air temperature at time t, $P_i^{HP}(t)$ is the HVAC power consumption at time t, θ_i denotes the building thermal inertia coefficient [20], ρ_i the HVAC efficiency factor and the sign $\pm$ reflects operation in heating (+) or cooling (−) mode. The HP operation is also subject to comfort constraint ensuring that the indoor temperature remains within acceptable bounds:

$$\underline{T_i^{in}} < T_i^{in}(t) < \overline{T_i^{in}} \quad (6\text{-}9)$$

The prosumer tracks the desired temperature setpoint to maximize the thermal comfort of occupants:

$$f_i^{HP} = \xi_{ac} \sum_{t_H}^{t_H+(H-1)\Delta t} (T_i^{in}(t) - T_i^{in^*})^2 \quad (6\text{-}10)$$

Where deviations from the desired setpoint temperature $T_i^{in^*}$ are penalised through the comfort cost term ξ_{ac}, which balances energy flexibility provision with user comfort preservation.

6.2.1.4. PV Generation

The maximum PV generation output is governed by the forecasted, time-varying solar irradiance profile $\zeta^{PV}(t)$in conjunction with the system's rated capacity $\bar{P}_i^{PV}$. The actual output $P_i^{PV}(t)$ may be curtailed when necessary, and is constrained as follows:

$$0 < P_i^{PV}(t) < \zeta^{PV}(t)\bar{P}_i^{PV}(t) \quad (6\text{-}11)$$

where $\zeta^{PV} \in [0,1]$ represents the normalized irradiance profile.

The objective here is to minimize the amount of clean power that is curtailed.

$$f_i^{PV} = \xi_{PV} \sum_{t_H}^{t_H+(H-1)\Delta t} (\zeta^{PV}(t)\bar{P}_i^{PV}(t) - P_i^{PV}(t))^2 \quad (6\text{-}12)$$

Additionally, the prosumer aims to maximize the utilization of available PV output by using it to charge the BS and EV, which is achieved by minimizing the following objective:

$$f_i^{util} = (P_i^{PV} + P_i^{BS} + P_i^{EV})^2 \quad (6\text{-}13)$$

6.2.2 Objective function

Each prosumer solves a MPO problem to determine the optimal operating trajectories of its DERs. The MPO simultaneously computes:

- The desired power setpoints $P_i^{d,*}(t)$ for each DER d, and
- The upward and downward flexibility margins $\delta_i^{d,*}$(t) that can be offered around these setpoints.

This enables the prosumer to plan its own energy management strategy while explicitly quantifying the flexibility that can be provided to the wider system. The first objective of the MPO is to maximise the flexibility that can be offered by the prosumer. By identifying the feasible upward and downward reserves around the desired setpoints, the framework enables active participation in demand-response schemes and peer-to-peer (P2P) trading platforms. This valorises flexibility as a service while ensuring that device-level operational constraints are respected. A second objective is to maximise the utilisation of PV generation. The optimisation seeks to absorb locally available PV power through battery storage and EV charging. This strategy improves local self-consumption, reduces grid dependency, and enhances the integration of renewable energy into the system. The third objective addresses the costs associated with providing flexibility. These include battery and EV degradation due to frequent cycling, as well as potential deviations from thermal comfort when HVAC units are used for demand-side flexibility. By explicitly penalising these costs in the optimisation, the MPO ensures that flexibility is provided in a sustainable manner that preserves asset lifetime and protects user well-being. Finally, the optimisation aims to maximise the net power injection into the grid. This means either maximising export when generation exceeds demand, or minimising import when loads dominate. In both cases, the prosumer contributes to a more balanced and efficient grid while simultaneously benefiting from economic incentives linked to energy exports and reduced consumption from the grid. The MPO formulation is:

$$\min_{P_i^{d,}(t),\delta_i^d} \sum_{t\in\mathcal{H}} \sum_{d\in\mathfrak{D}} -\delta_i^d + f_i^d\left(P_i^d\right) + f_i^{util}\left(P_i^d\right) - P_i^{total}(t) \quad (6\text{-}14)$$

$$\textbf{Subject to } \underline{P_i^d}(t) \le P_i^d(t) - \delta_i^d,\ P_i^d(t) + \delta_i^d \le \overline{P_i^d} \quad (6\text{-}15)$$

All device-specific state constraints for each DER $d \in \mathfrak{D}$ (6-16)

$$P_i^{total}(t) = \sum_{d\in\mathfrak{D}} P_i^d(t) - P_i^{fixed}(t) \quad (6\text{-}17)$$

$$\mu_1\left|P_i^d\right| \le \delta_i^d \le \mu_2\left|P_i^d\right| \quad (6\text{-}18)$$

where $\mathfrak{D} \subseteq$ {BS, EV, HVAC, PV} is the set of all DERs owned by prosumer i, $f_i^d\left(P_i^d\right)$ represents device-specific comfort costs (e.g. battery degradation, EV SoC tracking, HVAC discomfort); $f_i^{util}\left(P_i^d\right)$ denotes PV utilization penalties, encouraging local consumption rather than curtailment, δ_i^d captures the feasible upward and downward flexibility margins associated with device d. The optimisation incorporates upper and lower limits on the flexibility of each DER ($\mu_1 < \mu_2$). These constraints guarantee that the flexibility range remains realistic while respecting the physical and operational capabilities of each device. Since the formulation involves absolute values, which are inherently non-convex, an exact big-M reformulation is used. This reformulation introduces auxiliary non-negative variables and binary decision variables, thereby ensuring tractability of the optimisation problem. The reformulation can be expressed as:

$$P_i^d = P_i^{d+} - P_i^{d-}, \left|P_i^d\right| = P_i^{d+} + P_i^{d-}, z_i^d \in \{0,1\} \quad (6\text{-}19)$$

$$0 \leq P_i^{d+} \leq z_i^d \overline{P_i^d} \quad (6\text{-}20)$$

$$0 \leq P_i^{d-} \leq (1 - z_i^d)\left|\underline{P_i^d}\right| \quad (6\text{-}21)$$

This formulation is necessary for devices such as BS and EVs, where bidirectional charging and discharging must be modelled. As a result, the overall MPO becomes a Mixed-Integer Quadratic Program (MIQP). This problem class allows the simultaneous treatment of continuous operational variables (e.g. power flows, temperatures, states of charge) and discrete operational states (e.g. charge/discharge mode), thereby providing a rigorous yet computationally tractable framework for the optimal valorisation of DER flexibility. At each decision-making interval t, once the MPO is solved, the prosumer aggregates the solutions and flexibility margins across all its DERs to determine the total baseline power injection $\sum_{d \in \mathfrak{D}} P_i^d$ together with the corresponding feasible flexibility range that can be offered to the market. This range reflects the maximum flexibility potential of the prosumer, since it is computed from device-level costs, operational constraints, and the physical characteristics of the DERs. It should be emphasised that this flexibility envelope does not incorporate higher-level strategic behaviour, but instead represents the intrinsic technical capabilities of the assets under optimal operation.

6.3 Simulation results

The optimisation framework for the optimal valorisation of DER flexibility was evaluated through numerical simulations. The purpose of these experiments is to demonstrate the capability of the MPO to:

(i) derive feasible and realistic flexibility margins from heterogeneous DERs,
(ii) ensure efficient utilisation of local renewable generation, and
(iii) quantify the aggregated flexibility that can be offered to the market or P2P energy community platform.

The use case focuses on a single residential prosumer equipped with a portfolio of DERs, namely: a rooftop photovoltaic installation, a stationary battery storage system, an electric vehicle, and a heating, ventilation, and air conditioning unit. Simulations are conducted with an hourly resolution over a planning horizon of 24 hours. Forecasts of PV generation and outdoor temperature are assumed to be available at the beginning of each optimization horizon. The household's consumption and PV generation profiles are derived from real residential datasets collected in typical houses in Madeira Island, ensuring that the simulations reflect realistic operating conditions. Additional information regarding these households can be found in Section 5.2.

Figure 6-1 illustrates a representative example of the aggregate household electricity demand together with the corresponding PV generation profile for a typical day.

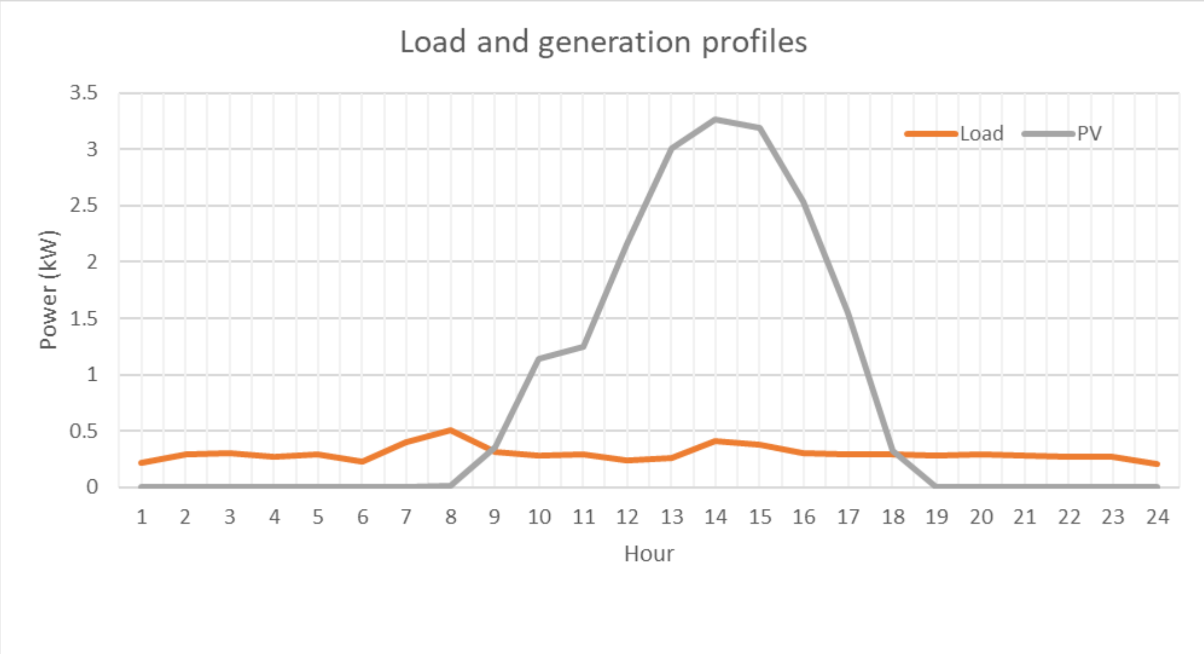

Figure 6-1: Representative household electricity demand and PV generation profile for a typical day

The test case is implemented in Python using the Pyomo optimisation framework, considering a 24-hour planning horizon with hourly resolution. The simulated prosumer is equipped with key DERs, including a 10-kWh battery storage system, a 6-kWh electric vehicle (EV) (available only during certain period of the day e.g., 13pm-22pm), and a HVAC system characterised by thermal dynamics parameters θ = 0.85 and ρ = 0.8. Indoor temperature comfort is maintained within [18°C, 24°C], with a setpoint of 21°C, and deviations are penalised (equation (6-8)). The outdoor temperature follows a realistic daily profile ranging from 5°C to 22°C. Table summarised the main technical parameters and operating constraints.

Table 6-1: Parameters of DERs in the Simulated Prosumer

DER Component	Parameter	Value / Description
Battery Storage System	Rated energy capacity	10 kWh
	Initial state of charge (SoC_0)	30%
	Efficiency	95%
	SoC limits	[10%, 90%]
Electric Vehicle	Rated energy capacity	6 kWh
	Availability window	13:00 – 22:00
	Initial SoC	30%
	Target SoC (by departure)	80%
HVAC System	Thermal dynamics parameters	θ = 0.85, ρ = 0.8
	Indoor temperature bounds	[18°C, 24°C]
	Comfort setpoint	21°C
	Outdoor temperature profile	5°C – 22°C (daily variation)

PV generation is modelled using forecasted hourly irradiance values, resulting in a bell-shaped production curve typical of residential systems. Residential load is forecasted with hourly granularity. These parameters establish a realistic simulation environment for assessing the flexibility of DERs under varying control strategies. To demonstrate the capabilities of the proposed framework, two representative scenarios are evaluated: (1) Baseline operation, in which DERs follow standard rule-based behaviour, and (2) Optimised operation, where DER scheduling is co-optimised to minimise costs while respecting comfort and operational constraints.

6.3.1 Baseline operation

In the baseline case, the DERs operate according to simple rule-based behaviour, without coordinated optimisation. The baseline operation assumes:

- The PV generation is directly used to cover the household's demand, and any surplus is exported to the grid.
- The stationary battery remains idle or follows a straightforward charging rule, storing excess PV generation until its maximum state of charge (SoC) is reached.
- The EV begins charging immediately when it becomes available and continues at maximum power until the required SoC is achieved by the departure time.
- The HVAC unit regulates indoor temperature around the setpoint by activating as needed, without providing flexibility to the system.

Figure 6-2 shows the typical energy flows in the baseline case, including PV generation, household demand, indoor and outdoor temperature, the HVAC power and the net grid exchange.

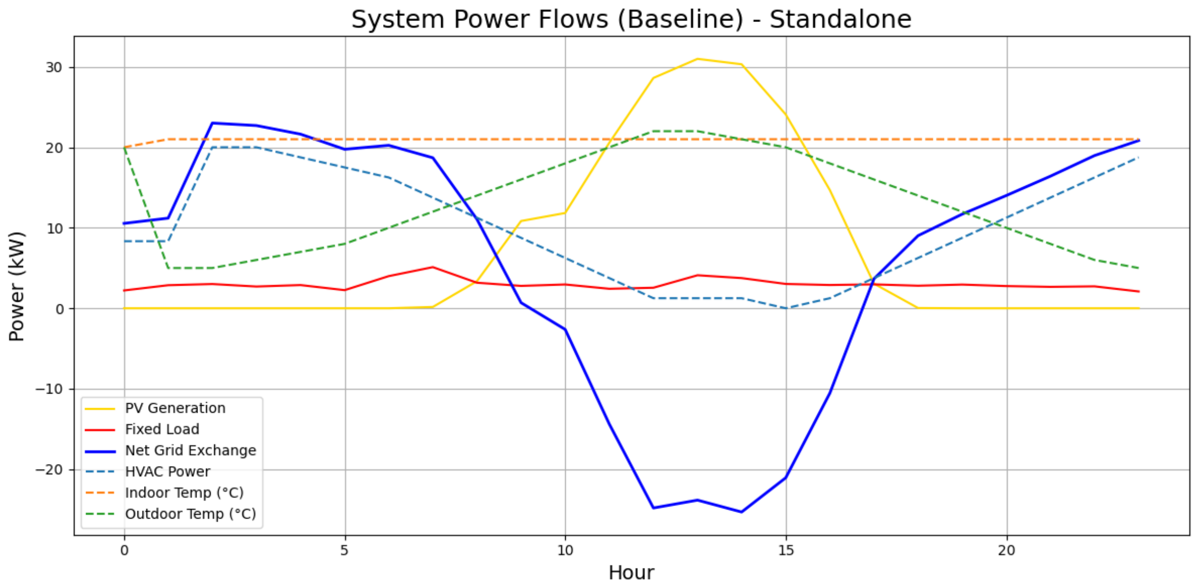

Figure 6-2: System power flows and temperature dynamics in the baseline scenario

The baseline simulation highlights the main characteristics and limitations of uncoordinated DERs operation. The PV output follows the expected diurnal profile, with generation peaking around midday. Since the stationary battery is idle, excess PV production cannot be stored and is directly exported to the grid. Conversely, during morning and evening hours when PV is unavailable, the household relies heavily on grid imports. The net grid profile shows significant variability. Imports dominate during low PV hours, while large surpluses are exported around noon. This behaviour indicates low self-consumption and limited valorisation of local renewable generation. The HVAC system maintains indoor temperature within the comfort bounds around the setpoint. However, HVAC consumption is not coordinated with PV availability, meaning additional electricity is often drawn from the grid during operation.

Overall, the baseline results demonstrate a lack of flexibility and coordination. The household fails to maximise local renewable utilisation, relies on imports during demand peaks, and does not provide explicit flexibility margins to the system. This underlines the necessity of optimisation strategies that coordinate DER operation for improved self-consumption, cost efficiency, and flexibility provision.

6.3.2 Optimized operation

In the optimised case, the DERs are scheduled in a coordinated manner through the flexibility optimisation framework. Unlike the baseline scenario, the objective is not only to satisfy local demand but also to maximise self-consumption, reduce grid dependency, and provide flexibility services. The PV generation is prioritised for self-consumption, with excess energy optimally stored in the stationary battery or used to supply the EV when available. The battery operates dynamically, charging during periods of surplus generation or low demand, and discharging when demand is high thereby reducing grid imports. The EV charging is shifted and modulated within the available connection window to meet the target state of charge (SoC) by departure time, while minimising peak demand and aligning charging with PV availability where possible. The HVAC unit is scheduled to maintain indoor temperature within comfort bounds while adjusting power consumption dynamically to provide demand flexibility.

Figure 6-3 presents the optimised operation results, including the evolution of PV generation, fixed load, HVAC power, indoor and outdoor temperature, the net grid exchange, and the coordinated behaviour of the battery and EV.

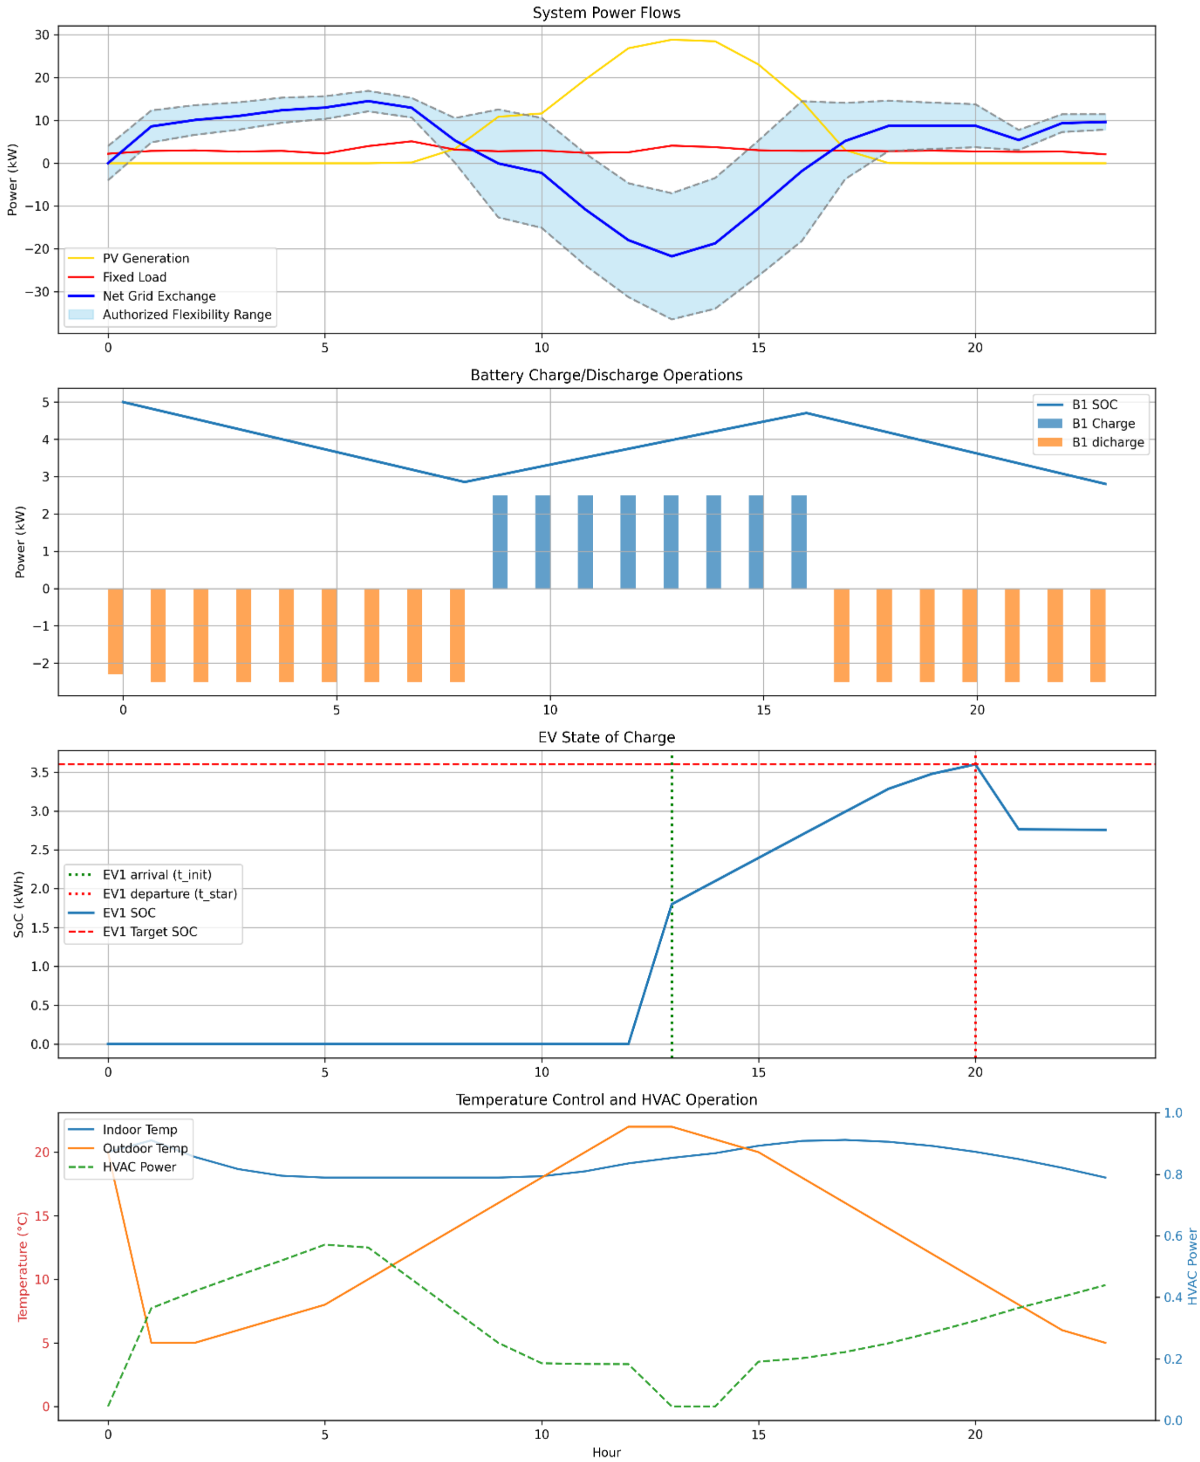

Figure 6-3: System power flows and temperature dynamics in the optimised scenario

The results presented in **Figure 6-3** highlight the benefits of coordinated operation of the DER portfolio compared to the baseline case. Several important behaviours can be observed. The authorised flexibility range shows that the household is now able to provide upward and downward flexibility across most of the day. This is enabled not only by PV curtailment when available but also by active participation of the battery and EV, which increases the system's capability to respond to external flexibility requests even in hours without solar production. The battery exhibits a balanced charge/discharge pattern, charging mainly during midday when PV generation is abundant, and discharging in the evening and early morning when demand is high. This behaviour reduces reliance on the grid and smooths the net demand profile. Unlike in the baseline scenario, the battery never remains idle and continuously supports the optimisation goals. The EV shifts its charging closer to periods of higher PV generation and lower demand, while ensuring the target SoC is reached before departure. This flexible charging profile avoids unnecessary grid peaks and leverages local renewable energy, demonstrating the optimisation's ability to align EV operation with both system and user objectives. The HVAC unit adapts its operation dynamically to maintain indoor comfort within the predefined bounds while also contributing to flexibility. Instead of following a rigid thermostat behaviour, its power consumption is modulated depending on system conditions, enabling temporary load reductions or increases without violating comfort constraints. This provides an additional source of flexibility beyond storage and EVs. Grid imports are significantly reduced compared to the baseline, particularly during PV production hours when demand is largely covered by local resources. The smoother net exchange profile indicates improved self-consumption and reduced stress on the grid, while also showcasing the household's ability to act as a flexible prosumer.

To further evaluate the performance of the proposed framework, several key performance indicators (KPIs) were computed, mainly those defined in Section 4.2.2.2, including the self-consumption ratio (SCR), self-sufficiency ratio (SSR), and peak power exchanges with the main grid. These metrics provide quantitative insight into how efficiently local renewable generation is utilized and how much the household depends on the external grid. Table 6-2 summarizes the KPI values obtained for both the baseline and optimized scenarios. The results show a clear improvement in energy autonomy and grid interaction efficiency when the optimization framework is applied, highlighting its effectiveness in enhancing self-consumption, reducing imports, and mitigating peak power flows.

Table 6-2: Comparison of Key Performance Indicators between Baseline and Optimised Operation

Indicator	Description	Baseline scenario	Optimised scenario
Self-Consumption Ratio (SCR)	Share of locally produced PV energy consumed on-site (Further details please see section 4.2.2.3)	0.31	0.51
Self-Sufficiency Ratio (SSR)	Share of total demand supplied by local generation (Further details please see section 4.2.2.4)	0.18	0.38
Peak Import Power (kW)	Maximum hourly grid import	23.01	14.49
Peak Export Power (kW)	Maximum hourly grid export	25.31	21.71

From the results presented in Table 6-2, we can see in the baseline operation, the self-consumption ratio (SCR) is 0.31, meaning that only about 31% of the locally generated energy is consumed on-site, while the rest is exported to the grid. The self-sufficiency ratio (SSR) reaches 0.18, indicating that only 18% of the household's energy needs are supplied by local generation, with the majority of demand covered by grid imports. Additionally, both the import and export peak powers are relatively high (≈23 kW and 25 kW), reflecting the lack of coordination between the distributed energy resources (DERs).

In contrast, under optimized operation, coordinated scheduling of the DER portfolio substantially enhances performance. The SCR increases to 0.51, and the SSR nearly doubles to 0.38, highlighting a much higher utilization of local renewable energy and reduced grid dependency. The optimization also reduces peak power exchanges import peaks drop by 37% (from 23.0 kW to 14.5 kW) and export peaks by 14% (from 25.3 kW to 21.7 kW) which contributes to mitigating stress on the distribution grid. Overall, the optimised coordination of PV, battery, EV, and HVAC results in improved energy autonomy, reduced grid dependency, and increased flexibility provision. This illustrates the effectiveness of the proposed optimisation framework in leveraging distributed resources to provide both economic and system-level benefits.

Overall, the optimised coordination of PV, battery, EV, and HVAC results in improved energy autonomy, reduced grid dependency, and increased flexibility provision. This illustrates the effectiveness of the proposed optimisation framework in leveraging distributed resources to provide both economic and system-level benefits.

7 Conclusions

Deliverable D2.2 presented the proof of concept of four complementary methods aimed at supporting consumers and ECs in defining, optimising, and valorising their operational strategies. The work carried out under Task 2.2 has demonstrated that advanced data-driven tools can effectively assist both individual actors and community managers in understanding their behaviour, anticipating system conditions, and identifying optimal participation strategies in energy sharing, flexibility, and market mechanisms.

The first key outcome concerns the development of Machine Learning based Forecasting methods by EIFER. The implemented models provide accurate and adaptable predictions for electricity consumption, renewable generation (PV and wind), electric vehicle demand, and market price signals. The forecasting results were shown to be robust across different time horizons and data granularities, enabling the anticipation of operational conditions and supporting downstream optimisation modules. The flexibility of the modelling approach allows future integration of new data sources or external signals, such as weather uncertainty and policy scenarios, which will be further exploited in Task 4.2 when the forecasting engine is embedded into the U2Demo platform.

The Data Analytics methods, developed by KU Leuven, constitute the second major result of this task. They provide quantitative evaluation criteria capable of describing and assessing the behaviour of consumers and ECs from technical, economic, and social perspectives. Technical indicators such as self-consumption and self-sufficiency ratios, peak power flows, and grid interaction parameters provide a detailed understanding of energy dynamics at both individual and community levels. Economic metrics quantify cost savings and efficiency, offering tangible incentives for participation. Social indicators like fairness indices ensure equitable benefit distribution, fostering trust and inclusiveness. Environmental metrics, particularly carbon reduction calculations, highlight the ecological benefits of local energy strategies. These data analytics indicators can be applied to reference baselines for consumption, production, and storage, which serve as a benchmark to evaluate the impact of optimisation or decision-making interventions. In later project phases, these same data analytic indicators will be used to monitor real demonstration results and to quantify the benefits of energy sharing and flexibility actions in practice.

The decision support methods, developed by INESC ID, address community-level operation and coordination of distributed energy resources. The proof-of-concept implementation within the PyECOM platform demonstrated how multi-objective optimisation can balance economic, environmental, and operational goals within an energy community. The results confirm that structured decision-support tools can provide actionable strategies for EC managers, helping them to define operating policies, that can enable flexibility and peer-to-peer trading. The modular design of these methods ensures scalability and compatibility with the forecasting and analytics components, preparing their integration within the forthcoming demonstrator environment.

The fourth method, the DER Valorisation methods developed by R&D Nester, focuses on peer-level decision making and optimisation of individual resources. The proposed framework performs multi-period optimisation of DERs, covering PV, batteries, EVs, and flexible loads, over a defined horizon. It ensures compliance with technical and comfort constraints while maximising self-consumption and flexibility value. The developed DER valorisation tool demonstrates the ability to effectively coordinate and optimise the operation of distributed energy resources within a residential prosumer setting. By integrating PV, stationary storage, EV charging, and HVAC flexibility into a unified optimisation framework, the tool enhances self-consumption, reduces reliance on the grid, and quantifies the latent flexibility potential of household assets. This enables prosumers to actively participate in emerging energy markets and contribute to system-level stability. Ultimately, the tool highlights the strategic value of DERs as active resources in the energy transition, bridging the gap between individual prosumer benefits and collective grid resilience, and creating additional value through participation in flexibility markets and P2P transactions.

A central achievement of Task 2.2 is the definition of the methodological framework that connects these four developments. Forecasting provides the predictive inputs required by the other methods. Data Analytics validates and characterises system behaviour, while both Decision Support and DER Valorisation generate optimised operational strategies based on those inputs. This framework defines not only the logical flow of information but also the feedback mechanisms necessary for continuous refinement and adaptation. Although each method has been developed and validated independently within this deliverable, their dependencies have been explicitly designed to ensure future interoperability. This guarantees that, in the next project phases, they can be seamlessly combined within the U2Demo platform to enable end-to-end data-driven decision support for consumers and ECs.

From a technical perspective, Task 2.2 successfully demonstrated the feasibility of implementing complex analytical and optimisation tools using realistic data, representative use cases, and consistent modelling assumptions. The results confirm the soundness of the methodological design and provide a solid starting point for scaling up the developments towards higher TRLs. In Task 4.2, each method will be further tested, refined, and integrated within the platform, ensuring operational compatibility, data exchange, and alignment with the demonstration objectives.

In addition to its direct contribution to Task 4.2, the outputs from Deliverable 2.2 also feed other activities of the U2Demo project. The forecasting, analytics, and optimisation results serve as inputs to Task 2.3, supporting the definition of energy sharing and P2P trading models, and can also be exploited in Task 2.4 for the development of mechanisms for market and flexibility participation, including system services for TSOs and DSOs. Hence, the methodologies described in this deliverable are not standalone tools but key building blocks that will underpin several subsequent developments within the project.

In conclusion, Deliverable 2.2 consolidates the methodological and technical groundwork required for enabling intelligent, data-driven decision-making at both peer and community levels. The demonstrated approaches collectively establish the foundation for the U2Demo platform, where predictive modelling, behavioural analytics, and optimisation will be combined to enhance energy sharing, flexibility utilisation, and market participation. The outcomes achieved under Task 2.2 validate the consistency of the overall approach and provide a clear pathway towards implementation and demonstration in the next phase of the project.

8 Bibliography

[1] Hyndman, R. J., & Khandakar, Y., “Automatic Time Series Forecasting: The forecast Package,” *R. Journal of Statistical Software,* pp. 27(3), 1–22., 2008.

[2] Rob Hyndman, Anne Koehler, Keith Ord, Ralph Snyder, Forecasting with Exponential Smoothing, Springer Berlin, Heidelberg, 2008.

[3] Assimakopoulos, V. , Nikolopoulos, K., “The theta model: a decomposition approach to forecasting,” *International Journal of Forecasting,* vol. 16, 2000.

[4] 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 Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17),* p. 3149–3157, 2017.

[5] David Salinas, Valentin Flunkert and Jan Gasthaus, “DeepAR: Probabilistic Forecasting with Autoregressive Recurrent Networks,” *Amazon Research,* 2019.

[6] Bryan Lim, Sercan O. Arik, Nicolas Loeff and Tomas Pfister, “Temporal Fusion Transformers for Interpretable Multi-horizon Time Series Forecasting,” *International Journal of Forecasting,* vol. 37, no. 4, pp. 1748-1764, 2020.

[7] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, Jayant Kalagnanam, “A Time Series is Worth 64 Words: Long-term Forecasting with Transformers,” *International Conference on Learning Representations,* 2023.

[8] Abhimanyu Das, Weihao Kong, Andrew Leach, Shaan K Mathur, Rajat Sen and Rose Yu, “Long-term Forecasting with Ti{DE}: Time-series Dense Encoder,” *Transactions on Machine Learning Research,* 2023.

[9] Abdul Fatir Ansari, Lorenzo Stella, Ali C. Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, O. Shchur, S. S. Rangapuram, Sebastian P. Arango, S. Kapoor, J. Zschiegner, D. C. Maddix, Hao W. , M. W. Mahoney, K. Torkkola, A. G. Wilson, M. Bohlke-Schneider, “Chronos: Learning the Language of Time Series,” *Transactions on Machine Learning Research,* 2024.

[10] “Open Meteo API,” [Online]. Available: https://open-meteo.com/.

[11] Hebrail, G. & Berard, A., “Individual Household Electric Power Consumption,” UCI Machine Learning Repository, 2006.

[12] DCC Public GIS Portal, 29 January 2025. [Online]. Available: https://data.dundeecity.gov.uk/search?collection=dataset&tags=charging%2520sessions%2Cev%2520charging.

[13] UK Power Networks, “Photovoltaic (PV) Solar Panel Energy Generation data,” 21 October 2025. [Online]. Available: http://data.europa.eu/88u/dataset/2nlqm.

[14] J. a. P.-F. P. J. a. G. M. C. Q. Sánchez-Soriano, “Historical Hourly Information of Four European Wind Farms for Wind Energy Forecasting and Maintenance,” *Data,* 2025.

[15] H. Fox, “European Wholesale Electricity Price Data,” Ember, 2025. [Online]. Available: https://ember-energy.org/data/european-wholesale-electricity-price-data/.

[16] E. Gomes, L. Pereira, A. Esteves, and H. Morais, “PyECOM: A Python tool for analyzing and simulating Energy Communities,,” *SoftwareX,* vol. vol. 24, p. p. 101580, Dec. 2023.

[17] Agência Portuguesa do Ambiente, “Fator de Emissao da Electricidade,” [Online]. Available:

https://apambiente.pt/sites/default/files/_Clima/Inventarios/FE_GEE_Electricidade_2024_final.pdf. [Accessed 03 Sep 2025].

[18] P. J. L. Amado, “Avaliação Energética de uma Unidade de Cogeração,” Instituto Politécnico do Porto (Portugal), Porto, 2016.

[19] R. E. N. (REN), “REN Data Hub – Electricity: Daily Balance,” [Online]. Available: https://datahub.ren.pt/pt/eletricidade/balanco-diario/. [Accessed 03 Sep. 2025].

[20] Vineet Jagadeesan Nair, Anuradha Annaswamy,, “A game-theoretic, market-based approach to extract flexibility from distributed energy resources,” *IFAC-PapersOnLine,* vol. 58, no. 30, pp. 163-168, 2024.