FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Mapping Trade-Offs for Aggregator Multi-Market Participation

João Pedro de Sousa Ribeiro



Mestrado em Engenharia Eletrotécnica e de Computadores

Supervisor: Doutor Tiago André Teixeira Soares

September 28, 2026

© João Pedro de Sousa Ribeiro, 2026

Resumo

A crescente integração de recursos energéticos distribuídos está a transformar o papel dos consumidores nos sistemas elétricos, criando novas oportunidades para a gestão ativa da procura, do armazenamento e da produção local. A agregação permite coordenar estes recursos e facilitar a participação em mecanismos de mercado que dificilmente estariam ao alcance de consumidores individuais. Contudo, a participação em diferentes mercados e serviços aumenta a complexidade operacional e introduz compromissos entre desempenho económico, requisitos técnicos e distribuição de benefícios.

Esta dissertação desenvolve e estende um sistema de gestão de energia agregado para coordenar múltiplos consumidores e avaliar o valor criado pela coordenação de recursos distribuídos e pelo acesso progressivo a diferentes mecanismos de mercado. O modelo é formulado como um problema determinístico de programação linear inteira mista e executado segundo uma abordagem de horizonte deslizante, preservando as restrições físicas de cada consumidor. Integra armazenamento em baterias, carregamento inteligente de veículos elétricos, flexibilidade da procura estimada a partir de dados históricos, mercado diário, reserva de restabelecimento da frequência com ativação manual (mFRR) e contratação bilateral de energia.

A metodologia é avaliada num caso de estudo residencial com cinco consumidores, utilizando dados reais de consumo e produção fotovoltaica de 2022. São consideradas duas configurações de referência e seis casos de otimização, nos quais novas funcionalidades são introduzidas sucessivamente. Esta estrutura permite quantificar o contributo incremental de cada mecanismo e analisar os resultados ao nível agregado e individual, bem como em diferentes condições sazonais.

Os resultados mostram que o valor económico resulta da combinação entre coordenação dos recursos e acesso a novas oportunidades de mercado. A otimização do armazenamento e do carregamento dos veículos elétricos, juntamente com a venda remunerada de energia e o acesso ao mercado diário, reduz o custo operacional variável anual em 47,9% face à referência totalmente passiva. A introdução da reserva mFRR representa o maior ganho incremental, elevando a poupança acumulada para 93,5%, sobretudo através da remuneração da capacidade disponível das baterias e das receitas associadas à ativação esperada. A flexibilidade da procura e a contratação bilateral conduzem a uma poupança total de 97,0% na configuração mais completa. A análise individual mostra ainda que os benefícios não são distribuídos uniformemente e que a maximização do valor económico pode divergir de objetivos como o autoconsumo e a autossuficiência.

O trabalho demonstra o potencial de um sistema de gestão de energia agregado para coordenar recursos distribuídos e explorar fontes complementares de valor. Mostra também que os benefícios obtidos dependem das condições de acesso aos mercados de reserva, das hipóteses determinísticas adotadas, da distribuição do valor entre participantes e da complexidade introduzida por cada nova camada de participação. Estes compromissos são essenciais para interpretar os resultados e orientar estratégias de agregação mais realistas e equilibradas.

Palavras-chave: sistemas de gestão de energia • agregação • flexibilidade da procura • mercados de energia e reserva

Abstract

The growing integration of distributed energy resources is transforming the role of consumers in power systems, creating new opportunities for the active management of demand, storage, and local generation. Aggregation coordinates these resources and facilitates participation in market mechanisms that are difficult for individual consumers to access. However, participation in different markets and services increases operational complexity and introduces trade-offs between economic performance, technical requirements, and the distribution of benefits.

This dissertation develops and extends an aggregated energy management system to coordinate multiple consumers and assess the value created through the coordination of distributed resources and progressive access to different market mechanisms. The model is formulated as a deterministic mixed-integer linear programming problem and implemented using a rolling-horizon approach, while preserving the physical constraints of each consumer. It integrates battery energy storage, smart electric vehicle charging, demand-side flexibility estimated from historical data, day-ahead market participation, manual Frequency Restoration Reserve (mFRR) provision, and bilateral energy procurement.

The methodology is evaluated through a residential case study comprising five consumers, using real 2022 consumption and photovoltaic generation data. Two reference configurations are considered alongside six optimisation configurations, in which new functionalities are introduced successively. This structure makes it possible to quantify the incremental contribution of each mechanism and analyse the results at aggregate and individual levels, as well as under different seasonal conditions.

The results show that economic value arises from the combination of resource coordination and access to new market opportunities. Optimised battery and electric vehicle scheduling, together with remunerated energy exports and access to the day-ahead market, reduce the annual variable operating cost by 47.9% relative to the fully passive reference. The introduction of mFRR reserve provision delivers the largest incremental gain, increasing cumulative savings to 93.5%, mainly through remuneration for available battery capacity and revenues associated with expected activation. Demand-side flexibility and bilateral energy procurement increase total savings to 97.0% in the most complete configuration. The individual analysis also shows that the benefits are not distributed uniformly among consumers and that maximising economic value may diverge from objectives such as self-consumption and self-sufficiency.

The work demonstrates the potential of an aggregated energy management system to coordinate distributed resources and exploit complementary sources of value. It also shows that the benefits achieved depend on the conditions governing access to reserve markets, the deterministic assumptions adopted, the distribution of value among participants, and the complexity introduced by each additional layer of participation. These trade-offs are essential for interpreting the results and guiding the development of more realistic and balanced aggregation strategies.

Keywords: energy management systems • aggregation • demand-side flexibility • energy and reserve markets

UN Sustainable Development Goals

The United Nations Sustainable Development Goals (SDGs) provide a global framework to achieve a better and more sustainable future for all. The framework comprises 17 goals[1] designed to address the world's most pressing challenges, including poverty, inequality, climate change, environmental degradation, peace, and justice. This dissertation is aligned with three of these goals through the development of an aggregated energy management system that coordinates distributed energy resources and enables their participation across multiple electricity market mechanisms.

The specific SDGs addressed have the following names:

SDG 7 Affordable and Clean Energy

SDG 9 Industry, Innovation and Infrastructure

SDG 12 Responsible Consumption and Production

SDG	Target	Contribution	Performance Indicators and Metrics
7	7.2	The aggregated energy management system coordinates the use of photovoltaic generation with battery storage, electric vehicle charging, and flexible demand, supporting the integration and use of locally generated renewable electricity.	Self-consumption ratio; self-sufficiency ratio; grid import and export volumes.
9	9.4	The multi-consumer optimisation framework supports more efficient use of existing distributed energy infrastructure while preserving individual technical constraints and enabling aggregated participation across multiple market mechanisms.	Battery charging and discharging; electric vehicle charging schedules; reserve availability.
12	12.2	Coordinated storage, demand shifting, and energy procurement support more efficient temporal use of electricity and existing distributed energy resources.	Shifted demand; battery charging and discharging; energy procurement mix.

[1] https://sdgs.un.org/goals

Recognitions

This work was carried out with the support of INESC TEC and supervised at this institution by Igor Roberto Rezende e Castro de Abreu.

Acknowledgements

I would like to begin by thanking everyone who has been part of my academic journey, and especially those who accompanied me throughout the development of this dissertation, which marks the end of my time at the Faculdade de Engenharia da Universidade do Porto.

I would like to thank my supervisor, Tiago Soares, for his availability, guidance, and support throughout the development of this dissertation.

I am also very grateful to Igor Rezende for all his support, availability, and time throughout this work, and for always being willing to discuss ideas, doubts, and possible solutions.

To my friends, thank you for making this journey lighter and easier, even during the moments of collective struggle that inevitably came along the way. Thank you for all the moments, conversations, laughs, and memories we shared throughout these years. I will always look back on this chapter of my life, and on all of you, with great affection.

To my girlfriend, Catarina, thank you for your unconditional support and for always believing in me and encouraging me in everything I do. Thank you for always lifting me up when things became difficult, and for all your understanding, love, and patience throughout this journey. Having you by my side made everything easier, and without you, this would have been much harder to achieve.

To my brother, Guilherme, thank you for making me want to become a better person every day and for giving me the motivation to be someone you can look up to. Without even realising it, you give me strength and remind me of the person I want to be.

To my mother, thank you for everything. Thank you for giving me all the support and opportunities I needed to follow an academic path that you would also have liked to pursue but did not have the chance to. Thank you for every effort and every sacrifice you have made, both throughout these years and throughout my entire life, to give me the opportunities and the life I have today. I would not be here without you. This achievement is more yours than it is mine.

Finally, to my father, who is not here to see this achievement. I am sure you would be incredibly proud. I hope that one day I can come close to being the incredible man and father you were.

João Ribeiro

"Ever tried. Ever failed. No matter. Try again. Fail again. Fail better."

Samuel Beckett

Contents

1 Introduction **1**
- 1.1 Context and Motivation . . . 1
- 1.2 Objectives . . . 4
- 1.3 Related Projects and Publications . . . 4
- 1.4 Structure of the Dissertation . . . 5

2 State of the Art **7**
- 2.1 Energy Management Systems for Consumers and Aggregators . . . 7
  - 2.1.1 Consumer- and Prosumer-Level Energy Management Systems . . . 8
  - 2.1.2 Aggregator-Level Energy Management Systems . . . 10
  - 2.1.3 Functional Architecture of Energy Management Systems . . . 11
  - 2.1.4 EMS Architectures and Coordination Structures . . . 13
- 2.2 Demand-Side Flexibility . . . 15
  - 2.2.1 Concept and Classification of Demand-Side Flexibility . . . 15
  - 2.2.2 Demand-Side Flexibility in B2B and B2C Contexts . . . 17
  - 2.2.3 Characterisation and Quantification of Demand-Side Flexibility . . . 18
  - 2.2.4 Integration and Aggregation of Demand-Side Flexibility . . . 20
- 2.3 Electricity Market Participation and Procurement Mechanisms . . . 21
  - 2.3.1 Day-Ahead Energy Markets and Price Signals . . . 22
  - 2.3.2 Ancillary Service Markets and Reserve Products . . . 24
  - 2.3.3 Power Purchase Agreements and Bilateral Energy Procurement . . . 27
- 2.4 Aggregation Frameworks and Business Models . . . 29
  - 2.4.1 Definition and Role of Energy Aggregators . . . 30
  - 2.4.2 Aggregation Paradigms and Organisational Models . . . 31
  - 2.4.3 Business Models, Revenue Streams and Market Coordination . . . 33
  - 2.4.4 Digital Platforms and Real-World Implementations . . . 35
- 2.5 Comparative Review and Research Positioning . . . 37
  - 2.5.1 Scope and Selection Criteria . . . 37
  - 2.5.2 Scientific Papers Review . . . 38
  - 2.5.3 Overview of Reviewed Contributions and Research Positioning . . . 40

3 Aggregated Energy Management System Model **43**
- 3.1 EMS Framework . . . 43
- 3.2 Battery Storage Modelling . . . 45
- 3.3 Flexible Load Modelling . . . 46
- 3.4 Electric Vehicle Modelling . . . 49
- 3.5 EMS Optimisation Model . . . 52
  - 3.5.1 Problem Description . . . 52

3.5.2 Mathematical Formulation . . . 53
3.6 Computational Implementation and Result Processing . . . 59

4 Case Study and Trade-Off Analysis **61**
4.1 Case Study Description and Simulation Setup . . . 61
4.1.1 Consumer Portfolio and Tariff Scenarios . . . 62
4.1.2 Distributed Energy Resources . . . 63
4.1.3 Energy and Reserve Market Price Inputs . . . 64
4.1.4 Computational Setup . . . 65
4.2 Progressive EMS Configurations . . . 66
4.3 Portfolio-Level Results . . . 67
4.3.1 Operating Cost Across Cases . . . 67
4.3.2 Energy Procurement and Selling . . . 68
4.3.3 BESS Scheduling and Reserve Provision . . . 70
4.3.4 Electric Vehicle Scheduling . . . 71
4.3.5 Demand-Side Flexibility Activation . . . 72
4.3.6 Bilateral Procurement . . . 74
4.4 Seasonal Analysis . . . 75
4.5 Consumer-Level Analysis . . . 77
4.6 Trade-Off Discussion . . . 80

5 Conclusion **83**
5.1 Conclusions . . . 83
5.2 Future Work . . . 85

References **87**

A Detailed Consumer-Level Results **99**

List of Figures

1.1 EU electricity generation mix by source, 2015-2025. 2
1.2 Structure of a modern electricity system. 3

2.1 Main components and interactions in a consumer- or prosumer-level EMS. 9
2.2 Simplified functional architecture of an EMS. 11
2.3 Demand profile reshaping through load shifting, peak shaving, and valley filling. 18
2.4 Aggregator-managed EMS interactions with resources and markets. 21
2.5 Temporal activation hierarchy of FCR, aFRR, mFRR and RR. 25
2.6 Simplified structure of a physical renewable PPA. 27
2.7 Coordination framework linking TSO, DSO, BRPs, aggregators, and DERs. . . . 29

3.1 General structure of the proposed EMS framework. 44
3.2 Flexible load modelling workflow. 47
3.3 Rolling-horizon treatment of an overnight EV session with fixed overlap decisions. 50

4.1 Waterfall decomposition of the annual variable operating cost across cases. . . . 67
4.2 Average hourly portfolio EV charging profiles. 71
4.3 Average hourly portfolio net load reshaping in Case E. 73
4.4 Daily portfolio operation under Case F: (a) summer day and (b) winter day. . . . 75
4.5 Annual variable cost per consumer under Baseline Simple and Case F. 77

List of Tables

2.1 Comparison of main EMS coordination structures discussed in the literature. . . . 14
2.2 Main classification dimensions of DSF. . . . 16
2.3 Main operational requirements for Portuguese mFRR participation. . . . 26
2.4 Main organisational paradigms for energy aggregation. . . . 33
2.5 Comparison of selected operational optimisation studies. . . . 40

4.1 BTN retail tariff levels ($\hat{\pi}_{i,t}^{tar}$) by time-of-use cycle. . . . 62
4.2 Contracted power ($\widehat{P}_{i}^{contr}$) and daily standing charge ($\hat{\pi}_{i}^{cp}$) per consumer. . . . 62
4.3 BESS parameters applied to all portfolio consumers. . . . 63
4.4 EV parameters per consumer. . . . 64
4.5 Overview of EMS configurations and active operational layers. . . . 66
4.6 Annual variable operating cost and savings relative to Baseline Simple across cases. 67
4.7 Annual portfolio grid exchanges from Baseline Simple to Case C. . . . 69
4.8 Decomposition of the annual variable cost before and after reserve provision. . . . 70
4.9 Portfolio energy purchase shares by procurement channel in Cases C and F. . . . 74
4.10 SSR and SCR per consumer under Baseline Simple and Case F. . . . 79
4.11 Principal benefit and limitation of each progressive EMS case. . . . 80

A.1 Share of energy purchased through each channel, per consumer, in Case C and Case F. . . . 99
A.2 Annual variable cost per consumer across all evaluated configurations. . . . 99
A.3 Reserve contribution and attributed revenue per consumer in Case F. . . . 100
A.4 SSR per consumer across all evaluated configurations. . . . 100
A.5 SCR per consumer across all evaluated configurations. . . . 100

List of Acronyms

ACER	European Union Agency for the Cooperation of Energy Regulators
aFRR	automatic Frequency Restoration Reserve
AMI	Advanced Metering Infrastructure
ARIMA	Autoregressive Integrated Moving Average
ATE	Alliance for Energy Transition
B2B	Business-to-Business
B2C	Business-to-Consumer
BEMS	Building Energy Management System
BESS	Battery Energy Storage System
BRP	Balance Responsible Party
BSP	Balancing Service Provider
BTN	Normal Low Voltage
CBC	COIN-OR Branch and Cut
DAM	Day-Ahead Market
DER	distributed energy resource
DERMS	Distributed Energy Resource Management System
DR	demand response
DSF	demand-side flexibility
DSM	demand-side management
DSO	Distribution System Operator
EEM	European Energy Market
EMS	Energy Management System
ERSE	Energy Services Regulatory Authority
EU	European Union
EV	Electric Vehicle
FCR	Frequency Containment Reserve
GHG	greenhouse gas
HEMS	Home Energy Management System
HVAC	Heating, Ventilation and Air Conditioning
INESC TEC	Institute for Systems and Computer Engineering, Technology and Science

IoT	Internet of Things
KPI	key performance indicator
MARI	Manually Activated Reserves Initiative
mFRR	manual Frequency Restoration Reserve
MIBEL	Iberian Electricity Market
MILP	Mixed-Integer Linear Programming
MIP	Mixed-Integer Programming
MPC	Model Predictive Control
OMIE	Iberian Electricity Market Operator
OMIP	Iberian Energy Derivatives Exchange
PICASSO	Platform for the International Coordination of Automated Frequency Restoration and Stable System Operation
PPA	Power Purchase Agreement
PRR	Recovery and Resilience Plan
PV	photovoltaic
REN	Redes Energéticas Nacionais
RES	renewable energy sources
RR	Replacement Reserve
SCADA	Supervisory Control and Data Acquisition
SCR	Self-Consumption Ratio
SDAC	Single Day-Ahead Coupling
SDG	Sustainable Development Goal
SEL	Smart Energy Lab
SOC	state of charge
SSR	Self-Sufficiency Ratio
TERRE	Trans European Replacement Reserves Exchange
TSO	Transmission System Operator
V2G	Vehicle-to-Grid
VPP	Virtual Power Plant

Chapter 1

Introduction

This chapter introduces the research presented in this dissertation. It begins by framing the role of energy flexibility in the transformation of electricity systems, with particular emphasis on distributed energy resource (DER) integration, demand-side flexibility (DSF), and aggregation mechanisms for Business-to-Business (B2B) and Business-to-Consumer (B2C) consumer portfolios. The chapter then presents the motivation for optimisation-based Energy Management System (EMS) frameworks capable of supporting aggregator multi-market participation and assessing the resulting operational and economic trade-offs. The main objectives of the work are subsequently defined, followed by the related research project and scientific publication associated with the dissertation. Finally, the organisation of the remaining chapters is presented.

1.1 Context and Motivation

The ongoing global energy transition is driven by the urgent need to reduce greenhouse gas (GHG) emissions and ensure the long-term sustainability of energy systems. International climate objectives, including the goal of reaching net-zero emissions by mid-century, require a profound transformation of the way energy is produced, delivered, and consumed [1]. At the European level, policy initiatives such as the European Green Deal and the "Fit for 55" package have reinforced this direction by promoting the large-scale deployment of renewable energy sources (RES), the electrification of end-use sectors, and improvements in energy efficiency [2], [3].

As a result, electricity systems are undergoing a structural transformation. Traditional centralised generation paradigms are increasingly complemented by decentralised architectures characterised by higher shares of variable and non-dispatchable renewable generation, particularly wind and photovoltaic (PV) resources. This evolution is visible in the electricity generation mix of the European Union (EU) between 2015 and 2025, during which wind and solar generation increased while fossil-based generation declined. In 2025, wind and solar accounted for a larger share of EU electricity generation than fossil fuels [4], as illustrated in Fig. 1.1.

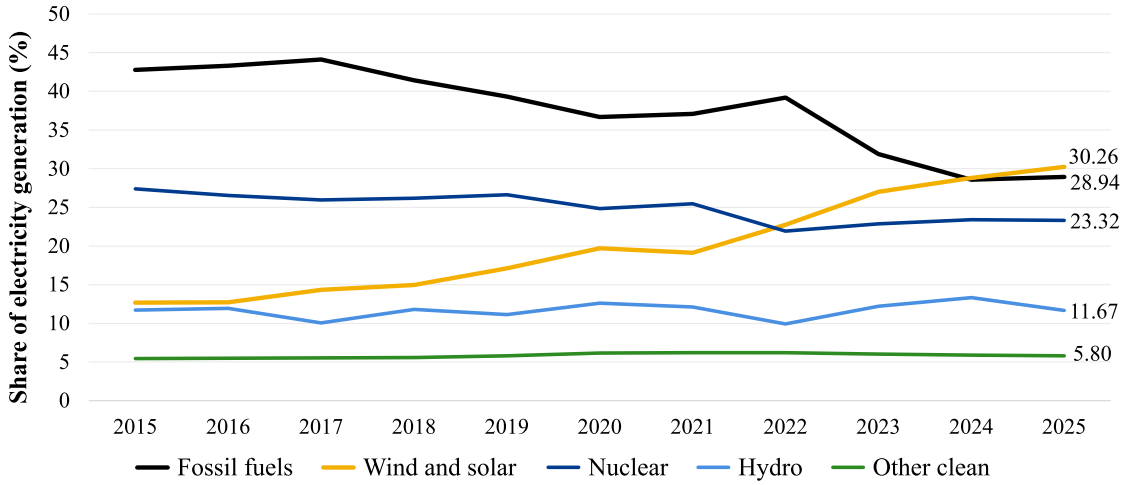


Figure 1.1: EU electricity generation mix by source, 2015-2025. Data source: [5].

Although the large-scale deployment of renewable generation is essential for decarbonisation, it also introduces important operational challenges. Variable renewable generation, particularly wind and solar, is weather-dependent and uncertain, increasing the variability of net load and placing additional pressure on system operation paradigms historically designed around centralised and dispatchable generation units [6]. This increases the need for alternative sources of controllability to maintain the real-time balance between electricity supply and demand. Recent energy price volatility and geopolitical disruptions have further exposed the vulnerability of energy systems to external supply shocks, reinforcing the need for more adaptive and resilient operational solutions [7].

In this context, flexibility has become a key resource for maintaining system reliability and economic efficiency. Broadly, flexibility refers to the capability of energy resources to adjust consumption, production, or storage operation in response to external signals, such as electricity prices, grid constraints, or activation requests from system operators. While flexibility has traditionally been provided mainly by conventional generation units, the decarbonisation of the power sector increases the importance of alternative sources, including energy storage, network-level operational solutions, and distributed flexibility originating from the demand side [8].

Among these alternatives, DSF is particularly relevant because it can support system operation without relying solely on additional generation capacity. It enables consumers to intentionally increase, decrease, or shift electricity consumption in response to economic incentives or control signals [9], contributing to peak shaving, congestion management, reserve provision, and the integration of high shares of variable renewable generation [10]. Fig. 1.2 illustrates the structure of a modern electricity system shaped by decarbonisation, decentralisation, and digitalisation, commonly referred to as the three Ds of the energy transition. It also shows where flexibility-enabling technologies fit within this structure.

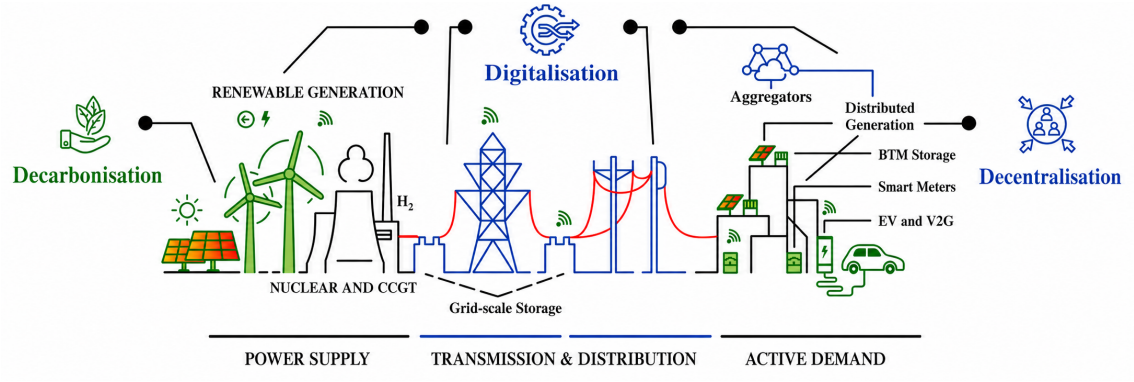


Figure 1.2: Structure of a modern electricity system. CCGT: combined-cycle gas turbine; BTM: behind-the-meter; EV: electric vehicle; V2G: vehicle-to-grid. Adapted from [11].

The relevance of DSF spans multiple layers of the electricity system, from local distribution networks to wholesale electricity markets and ancillary service mechanisms. In these contexts, flexible demand can improve operational efficiency, reduce the need for costly network reinforcements, and create new opportunities for active consumers and prosumers [10]. However, its effective use remains challenging because electricity demand is influenced by consumer behaviour, operational routines, and comfort requirements. Flexibility must also be predictable, measurable, reliable, and verifiable to be economically valued within market-based frameworks [6].

These challenges limit the direct participation of individual consumers in electricity markets and ancillary service mechanisms. Small-scale resources often lack the capacity, observability, and predictability required to satisfy minimum bid sizes, qualification rules, and performance guarantees [6]. European market designs and pilot projects have increasingly recognised the need to unlock distributed flexibility, but practical participation still depends on adequate coordination mechanisms and market-compatible representations of flexible resources [12].

Advanced decision-support tools and aggregation mechanisms are therefore essential to coordinate heterogeneous resources under technical, economic, and market constraints. EMSs provide an optimisation-based framework for scheduling DERs, storage systems, flexible demand, and controllable loads across different time horizons and market contexts [13], [14]. At the same time, aggregators can bundle flexibility from multiple consumers into market-compatible products while managing operational, commercial, and regulatory constraints [15], [16].

Despite these advances, existing EMS frameworks often address flexibility resources, market participation mechanisms, and procurement strategies separately, making it difficult to coordinate heterogeneous consumer portfolios while preserving individual operational constraints and assessing the resulting trade-offs.

From an aggregator perspective, local resources such as PV generation, a Battery Energy Storage System (BESS), an Electric Vehicle (EV), and flexible loads may support energy cost reduction, surplus selling, and reserve provision, while their operation must also be coordinated with bilateral procurement decisions. These options cannot always be maximised simultaneously, creating trade-offs between self-consumption, storage scheduling, electric mobility requirements,

demand reshaping, market exposure, reserve availability, and contractual procurement. These competing objectives motivate integrated EMS frameworks capable of coordinating heterogeneous distributed resources across multiple market layers while systematically evaluating the trade-offs associated with different scheduling and market participation strategies.

1.2 Objectives

This dissertation aims to develop an aggregated EMS for B2B and B2C consumer portfolios by extending a pre-existing single-site optimisation tool developed at the Institute for Systems and Computer Engineering, Technology and Science (INESC TEC). The framework focuses on the coordination of PV generation, BESSs, EVs, and flexible loads across energy markets, reserve provision, and bilateral procurement. The main goal is to assess how these resources can be coordinated within an aggregator-managed portfolio and to identify the operational and economic trade-offs that emerge as additional operational and market layers are progressively incorporated into the scheduling framework. The model was designed for this broader consumer scope but is evaluated through a residential B2C case study.

The specific objectives of this dissertation are:

- To review EMS-based approaches for consumers and aggregators, with emphasis on DSF, DERs, electricity market participation, reserve provision, and bilateral procurement.
- To formulate and implement an aggregated Mixed-Integer Linear Programming (MILP) model for rolling-horizon portfolio coordination that preserves the physical feasibility of each consumer.
- To integrate a data-driven DSF modelling layer based on historical consumption data, defining time-dependent upward and downward flexibility limits for each consumer.
- To integrate consumer-level BESS and EV operation, energy-market interaction, portfolio-level reserve provision and bilateral procurement within a common scheduling framework.
- To evaluate the implemented EMS through successive configurations using Smart Energy Lab (SEL) data, assessing the incremental impact of each modelling layer on portfolio and individual-consumer operation, economic performance, resource scheduling, and the trade-offs associated with aggregator multi-market participation.

1.3 Related Projects and Publications

The work developed within the scope of this dissertation is directly related to the objectives and activities of one research project, namely:

- Alliance for Energy Transition (ATE), a national mobilising agenda funded through Portugal's Recovery and Resilience Plan (PRR), focused on accelerating the decarbonisation of

the Portuguese energy sector through the integration of RES, the digitalisation of energy systems, and the development of innovative solutions across the energy value chain.

This dissertation resulted in one scientific paper presented at the 22nd International Conference on the European Energy Market (EEM). The corresponding publication is listed below:

- João P. Sousa, Igor Rezende and Tiago Soares, "Decision Framework for EMS-Controlled Flexible Loads and BESS: Optimal Market Selection for Energy and Ancillary Services", *2026 22nd International Conference on the European Energy Market (EEM)*, Trondheim, Norway, 2026, pp. 1-6, doi: 10.1109/EEM68581.2026.11589663.

1.4 Structure of the Dissertation

The dissertation is organised into five chapters, following a progressive structure from the general research context and state of the art to the proposed EMS formulation, simulation results, and final conclusions:

- Chapter 1 introduces the research topic, frames the motivation of the dissertation, defines the main objectives, and identifies the related project and publication.
- Chapter 2 reviews the main concepts related to EMSs, DSF, DERs, aggregation, electricity markets, reserve services, and bilateral procurement.
- Chapter 3 presents the aggregated EMS model developed and extended in this work, including the resource models, mathematical formulation, rolling-horizon operation, and computational implementation.
- Chapter 4 describes the case study and simulation setup, presents the evaluated EMS configurations, and analyses the operational and economic results at portfolio and individual-consumer levels and under different seasonal conditions, concluding with a synthesis of the resulting trade-offs.
- Chapter 5 summarises the main findings and contributions of the dissertation, discusses the limitations of the developed approach, and identifies possible directions for future work.

Chapter 2

State of the Art

This chapter reviews the concepts and research developments that support the proposed aggregated EMS. It first discusses EMSs at consumer, prosumer, aggregator, and portfolio levels, focusing on their functional architectures, coordination structures, and role in managing distributed resources. It then addresses DSF, including its classification, quantification, sector-specific characteristics, and aggregation.

The chapter also examines the market and procurement mechanisms most relevant to this dissertation, namely participation in the Day-Ahead Market (DAM), reserve products with emphasis on manual Frequency Restoration Reserve (mFRR), and bilateral procurement. These mechanisms shape aggregator decisions and create trade-offs between energy costs, reserve provision, flexibility use, and procurement strategies. Aggregation frameworks and business models are subsequently reviewed, before the chapter concludes with a comparative analysis of selected optimisation-based studies to identify research gaps and position the work developed in this dissertation.

2.1 Energy Management Systems for Consumers and Aggregators

EMSs have become increasingly relevant in electricity systems with growing shares of DERs. PV generation, storage systems, EVs, flexible loads, and other distributed assets are changing the way electricity is consumed, produced, stored, and exchanged. Managing these resources requires monitoring and coordination capabilities and, when multiple consumers or sites are involved, aggregation mechanisms that allow local assets to be operated according to technical constraints, economic signals, and operational objectives [17].

This coordination problem becomes more complex when controllable and partially controllable resources are considered together. Storage systems, electric mobility, local generation, and flexible demand may provide operational value, but their scheduling is constrained by availability, comfort requirements, state of charge (SOC) limits, uncertainty, and market conditions. In the Virtual Power Plant (VPP) literature, resource scheduling is therefore treated as a central energy-management task when distributed resources are coordinated for market-oriented operation [18].

The literature commonly distinguishes EMS applications according to their coordination scale. At consumer or prosumer level, EMSs typically manage local assets to reduce energy costs, increase self-consumption, limit peak demand, or preserve comfort. At aggregator or portfolio level, local capabilities can be coordinated and represented externally as energy, flexibility, or system-service products. VPPs and aggregator-level EMSs capture this broader coordination layer, where distributed resources remain physically dispersed but are operated through coordinated portfolio-level decisions [19].

This multi-level perspective introduces different technical and organisational requirements. VPPs and virtual aggregators are commonly described as coordinated associations of DERs that can create value for prosumers, support participation in energy markets, and provide services to system operators [20]. Their implementation depends not only on optimisation models, but also on communication infrastructures, forecasting tools, control strategies, market interfaces, and mechanisms for service provision. For this reason, centralised, decentralised, distributed, and hierarchical coordination structures are often discussed as alternative ways of organising EMS, VPP, and aggregator operation [20]. The following subsections review these aspects progressively, starting from consumer- and prosumer-level EMSs before moving towards aggregator-level operation, functional architectures, and coordination structures.

2.1.1 Consumer- and Prosumer-Level Energy Management Systems

Consumer- and prosumer-level EMSs operate at the interface between end-user behaviour, behind-the-meter assets, and external energy signals. In residential applications, a Home Energy Management System (HEMS) typically combines sensing devices, smart meters, user interfaces, controllable appliances, local generation, storage units, and communication with external actors to support automated scheduling decisions [21]. In building applications, a Building Energy Management System (BEMS) extends this logic to a broader set of subsystems, including Heating, Ventilation and Air Conditioning (HVAC), lighting, appliances, battery storage, and thermal storage. These systems commonly rely on strategies such as Model Predictive Control (MPC), demand-side management (DSM), optimisation, and fault detection to improve energy efficiency while preserving indoor comfort [22].

At the household level, demand controllability is shaped by the technical and behavioural characteristics of individual loads. Some appliances are essentially uncontrollable, whereas others can be curtailed, shifted in time, or operated within comfort-dependent limits. This distinction is important because the available flexibility of a local EMS is not determined solely by installed capacity, but also by appliance constraints, user preferences, and the acceptable delay or discomfort associated with each control action. Consequently, local scheduling objectives may extend beyond cost minimisation to include the peak-to-average ratio, user comfort, and operational feasibility, as illustrated by multi-objective approaches that explicitly represent economic, technical, and discomfort-related criteria [23]. A simplified representation of the main components and interactions typically involved in a consumer- or prosumer-level EMS is shown in Fig. 2.1.

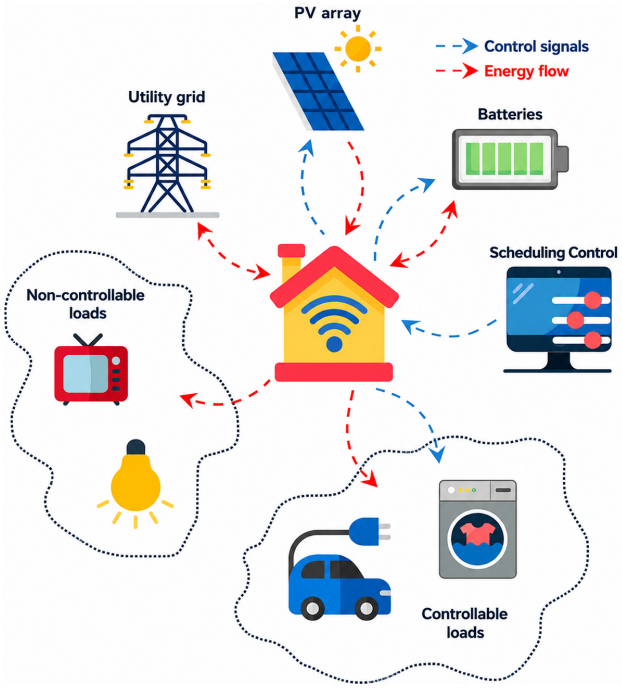


Figure 2.1: Main components and interactions in a consumer- or prosumer-level EMS [24].

The prosumer context introduces an additional layer of complexity, since local generation and storage modify the traditional role of the consumer. Rooftop PV generation, BESSs, and EVs allow households to shift consumption, store surplus renewable energy, and reduce imports during expensive periods. Double-layer HEMS strategies have been proposed to separate appliance scheduling from the dispatch of storage and EVs, improving the match between PV production and demand while keeping the optimisation problem more manageable [25]. Beyond day-to-day operation, EMS-based scheduling can also affect investment decisions, since the expected ability to shift demand and increase self-consumption influences the optimal sizing and economic value of PV-BESS systems [26].

Uncertainty remains a central challenge in consumer- and prosumer-level EMSs. Forecast errors in PV generation, dynamic electricity prices, appliance usage, and EV availability can reduce the value of schedules derived under deterministic assumptions. For this reason, robust and stochastic formulations have been proposed to improve performance under uncertain operating conditions. For instance, two-stage robust optimisation has been applied to smart-home energy management with PV, battery storage, EV charging, and demand response (DR), aiming to preserve cost savings under adverse renewable generation scenarios [27]. Such approaches can reduce sensitivity to forecast errors, although they also increase modelling and computational complexity.

Despite their benefits, isolated consumer-level EMSs still present important limitations when deployed at scale. Local optimisation can reduce individual bills and improve self-consumption, but uncoordinated responses to common price signals may also create synchronised demand shifts and new stress periods in distribution networks. HEMS surveys therefore emphasise the importance of communication architectures, coordination with utilities, and demand-response integra-

tion for smart-grid deployment [28]. Recent residential EMS reviews also indicate that practical deployment still requires further attention to system-level constraints, communication requirements, user-centred behaviour, and network-aware operation [29]. These limitations motivate the transition from isolated local optimisation towards aggregation mechanisms capable of coordinating distributed flexibility while preserving the technical and behavioural constraints of each consumer.

2.1.2 Aggregator-Level Energy Management Systems

Aggregator-level EMSs extend energy management beyond individual consumers by coordinating multiple distributed resources as a portfolio. This role is particularly relevant because the flexibility available at household, building, or service-sector level is often too small, dispersed, or heterogeneous to interact directly with electricity markets and system operators. By pooling controllable loads, distributed generation, storage systems, and EVs, aggregators can transform fragmented local flexibility into a more reliable, market-compatible, and operationally usable portfolio capability [30].

From an operational perspective, aggregation is not limited to combining individual schedules. Resource aggregators must identify and qualify flexible resources, forecast relevant information, define bidding or trading strategies, activate contracted flexibility, monitor delivery, and settle the delivered response [31]. These tasks require continuous information exchange between consumers, the aggregator, and external entities, since market opportunities must be evaluated together with physical constraints, user preferences, and the response characteristics of the participating assets.

The portfolio perspective also changes the way EMS models are formulated. Instead of optimising a single household or building, an aggregator must coordinate heterogeneous resources while preserving the feasibility of each participant. Recent studies on renewable energy communities illustrate this logic through centralised EMS formulations in which an aggregator computes reference power exchange profiles for local EMSs, allowing community-level objectives to be pursued while individual users retain their own operational behaviour [32]. This type of structure is particularly relevant when local and collective objectives are not perfectly aligned, since participants may behave cooperatively, partially cooperatively, or according to their own economic interests.

A further modelling challenge concerns the interaction between different types of flexibility. Aggregated portfolios may include PV generation, BESSs, EVs, thermal systems, power-to-heat units, and time-shiftable loads, each with specific temporal constraints and technical dynamics. Holistic MILP formulations have been proposed to represent these interactions in day-ahead operation, showing that the value of aggregation depends not only on the availability of individual assets, but also on the synergies between resource scheduling and trading options across wholesale, local, and internal trading layers [33]. However, this level of modelling detail increases computational complexity and may limit scalability when large portfolios are represented with high technical resolution.

Implementation aspects are equally important, especially when flexibility must be activated under market or system-service requirements. In industrial DR, multi-agent system architectures have been proposed to connect VPP operators with flexibility providers, supporting scheduling, activation, monitoring, and accounting processes through standardised communication sequences and interfaces [34]. Such approaches show that aggregator-level EMSs are not only optimisation models, but also operational platforms that must handle communication delays, interface heterogeneity, automation levels, and regulatory requirements. This motivates a closer look at the functional architecture of EMSs, where data acquisition, forecasting, optimisation, control, and external communication are organised into an integrated decision-making framework.

2.1.3 Functional Architecture of Energy Management Systems

Beyond their specific application domain, EMSs can be viewed as modular platforms that convert measurements, forecasts, and external signals into schedules, set-points, or control actions. This functional perspective separates the main stages of information processing from the particular assets, consumers, or markets being managed. In general terms, these stages include data acquisition and monitoring, forecasting and data processing, optimisation and decision-making, scheduling and control, and interfaces with users, assets, markets, or system operators [35]. This structure is especially relevant in aggregated applications, where information from multiple assets and consumers must be processed before coordinated decisions can be issued. Fig. 2.2 provides a generic representation of these functional interactions.

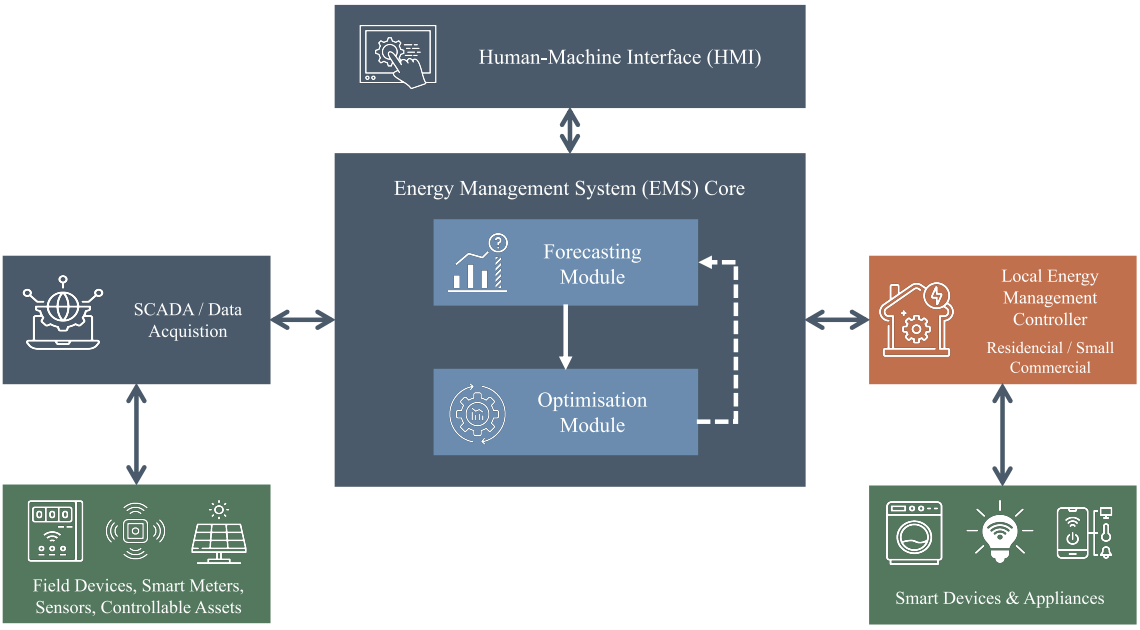


Figure 2.2: Simplified functional architecture of an EMS.

The starting point of this architecture is observability. Measurements from smart meters, sensors, field devices, controllable assets, and communication gateways provide the information required to characterise the current operating state of the system. At the residential scale, HEMS architectures typically include monitoring, logging, management, control, and alarm functionalities, supported by communication between the central controller, appliances, sensors, energy resources, and the user interface [36]. At higher levels, similar principles are extended through Supervisory Control and Data Acquisition (SCADA) systems, Advanced Metering Infrastructure (AMI), smart inverter measurements, and network models, which support real-time awareness in Distributed Energy Resource Management System (DERMS) platforms [37]. To strengthen this real-time observability, state-estimation functions can combine real-time measurements with accurate network models to determine current grid conditions, improve unreliable measurements, and identify potential violations in the distribution grid [37].

Raw measurements become operationally useful once they are processed into forecasts, indicators, or flexibility representations. Forecasting modules estimate variables such as demand, renewable generation, electricity prices, and resource availability, while data analytics modules extract patterns related to consumption behaviour, asset operation, or flexibility potential. In Internet of Things (IoT)-based implementations, data collection from renewable generation and smart devices can be combined with machine-learning forecasting and load-priority logic to support day-ahead or short-term scheduling decisions [38]. However, the increasing value of smart-meter and HEMS data also raises privacy and data-governance concerns. For this reason, some architectures favour local forecasting or the exchange of models and flexibility representations instead of raw personal data [39].

At the decision-making stage, forecasts, constraints, preferences, and external signals are translated into operational decisions. Depending on the problem structure, this stage may rely on linear or mixed-integer programming, MPC, stochastic or robust optimisation, metaheuristics, artificial intelligence methods, or hybrid approaches. Reviews of microgrid EMSs show that mixed-integer formulations are widely used because they can represent discrete operating modes, storage dynamics, scheduling constraints, and intertemporal decisions, although their computational burden can increase with the number of assets, binary variables, and uncertainty scenarios [40]. Therefore, the choice of optimisation method is closely linked to the EMS scale, time horizon, modelling detail, and required execution speed.

The final stage closes the loop between monitoring and action by communicating schedules, set-points, or control signals to the relevant devices and stakeholders. In local systems, this may involve appliance automation, user notifications, smart plugs, or energy management controllers. In aggregated or grid-facing systems, it may involve exchanging flexibility offers, market bids, operating limits, or control signals with an aggregator, a Distribution System Operator (DSO), a Transmission System Operator (TSO), or an external platform [37]. This implementation layer introduces practical challenges related to interoperability, communication delays, cybersecurity, and the division of control authority. Data ownership and privacy concerns are also relevant when smart-meter or HEMS information is exchanged beyond the consumer premises [39].

2.1.4 EMS Architectures and Coordination Structures

The previous subsection described the functional layers through which EMSs transform measurements, forecasts, and external signals into operational decisions. A complementary perspective concerns how these functions are coordinated. In practice, the architecture of an EMS defines where decisions are made, how much information is exchanged, and how control authority is distributed among local devices, consumers, aggregators, and system-level operators. The literature commonly distinguishes between centralised, decentralised, distributed, and hierarchical coordination structures, each reflecting a different compromise between global optimisation, local autonomy, communication burden, privacy, and implementation complexity [41].

In a centralised architecture, measurements and forecasts from local resources are sent to a supervisory controller, which computes schedules or set-points for the whole system [41]. This structure is attractive when the EMS has reliable access to system-wide information, since the optimisation problem can explicitly consider shared constraints, collective objectives, and external signals. However, the same dependence on a single decision point creates important weaknesses. As the number of assets increases, the computational burden and communication requirements grow, while failures or cyber-attacks affecting the central controller can compromise the entire operation. For this reason, the scalability and resilience of centralised EMSs become increasingly challenging as portfolios grow and include heterogeneous, independently owned resources [42].

Decentralised and distributed structures shift part of this decision-making capacity towards local controllers. In a decentralised EMS, each participant or subsystem optimises its own operation using local information and limited external signals. In a distributed EMS, local controllers also exchange selected information with neighbouring agents or peers to coordinate their decisions. This can improve scalability and privacy, because complete device-level profiles do not need to be transmitted to a central platform, and local controllers can react more quickly to disturbances. In networked microgrids, peer-to-peer and event-triggered communication schemes have been discussed as ways to support coordination while reducing unnecessary data exchange [43]. Similar principles also appear in VPP and aggregator platforms, where geographically dispersed DERs may be coordinated through decentralised, distributed, or comprehensive control arrangements [20]. The drawback is that local autonomy can make it harder to guarantee a globally optimal or network-feasible outcome, especially when participants have different objectives, incomplete information, or weak communication links.

Hierarchical EMSs attempt to combine these two perspectives by separating decision-making across levels and time scales. Fast local controllers preserve asset-level or microgrid-level feasibility, while upper layers solve slower coordination problems such as power exchange, market scheduling, or portfolio-level dispatch. In networked microgrids, this logic is often represented through lower-level EMSs that optimise individual microgrids and communicate surplus or shortage information, storage status, and marginal cost signals to an upper-level EMS, which then coordinates inter-microgrid power exchange and sends reference decisions back to the local level [44]. Applied two-level MPC structures illustrate the same principle. A supervisory layer computes ag-

gregate power and demand-shifting references, while local controllers track those references using their own generation, storage, and demand information [45]. The value of hierarchy lies less in removing complexity than in organising it into layers with different response times, information requirements, and control responsibilities. Table 2.1 summarises the main coordination structures discussed in the literature.

Table 2.1: Comparison of main EMS coordination structures discussed in the literature.

Structure	**Decision logic**	**Main advantages**	**Main limitations**
Centralised	A supervisory controller computes system-level decisions using system-wide information.	Global visibility, coordinated dispatch, and explicit treatment of shared constraints and collective objectives.	Scalability limitations, high communication requirements, and dependence on a single decision point.
Decentralised	Local controllers make decisions mainly from local information and limited external signals.	Local autonomy, reduced data-sharing requirements, and fast response to local operating conditions.	Limited global coordination and potentially suboptimal system-level operation.
Distributed	Local controllers coordinate through selected information exchange with neighbouring agents or peers.	Improved scalability, reduced central computation, and support for peer-to-peer coordination.	Dependence on reliable communication, together with cybersecurity and convergence challenges.
Hierarchical	Local and supervisory layers coordinate decisions across different time scales.	A balance between local responsiveness and global coordination, with a clearer allocation of control responsibilities.	Higher implementation complexity and the need for synchronisation and information exchange between layers.

The choice of EMS architecture is therefore not purely algorithmic. It depends on the scale of the system, ownership structure, communication infrastructure, cybersecurity requirements, market participation needs, and the degree of autonomy granted to end-users [41]. A highly centralised design may provide stronger global coordination, but it can require extensive data exchange and become difficult to scale. A highly decentralised design may preserve autonomy and privacy, but it can struggle to deliver system-level guarantees. Hierarchical approaches can provide a practical compromise by organising decision-making across local and supervisory layers, preserving local feasibility while enabling higher-level coordination [44].

This architectural choice is particularly important for DSF. The same flexible load may be treated as a direct control variable, a locally optimised preference-constrained action, or an aggregated flexibility envelope, depending on the EMS structure. This distinction becomes especially relevant when flexible resources from multiple B2B and B2C consumers must be coordinated without ignoring individual technical and behavioural constraints. The next section therefore shifts the focus from how EMSs are organised to the flexible demand-side resources that these architectures aim to monitor, schedule, and activate.

2.2 Demand-Side Flexibility

DSF has become increasingly important as power systems integrate larger shares of variable renewable generation and experience growing electrification of end-use sectors. In this context, flexibility must be provided not only by conventional generation and storage resources, but also by consumers capable of adjusting their electricity demand in response to system conditions. Demand-side flexibility can be understood as the capability of demand-side resources to reduce, increase, or shift electricity consumption over a specific period to support renewable integration, reduce peak demand, and lower system operating costs [46].

From an operational perspective, DSF turns consumption into a resource that can be characterised, scheduled, activated, and verified. Its value depends not only on the amount of power that can be modified, but also on attributes such as duration, response time, location, controllability, predictability, and the purpose of the flexibility service [47]. These attributes determine whether flexible demand can support local energy management, portfolio optimisation, congestion management, reserve provision, or participation in electricity markets. The following subsections review the concept and classification of DSF, its specific characteristics in B2B and B2C contexts, methods for characterising and quantifying flexibility, and the aggregation mechanisms required to mobilise distributed flexibility at scale.

2.2.1 Concept and Classification of Demand-Side Flexibility

Although the terms DSF, DSM, and DR are often used in related contexts, they should not be treated as exact synonyms. DSM refers to the broader set of strategies used to influence electricity consumption patterns, whereas DR refers more specifically to temporary changes in electricity use in response to prices, incentives, or system needs. DSF is the underlying capability that enables these actions. It represents the technical and behavioural ability of demand-side resources, including controllable loads and behind-the-meter assets, to modify their net electricity demand within acceptable constraints [48]. This distinction is important because not all available flexibility is directly dispatchable, measurable, or suitable for market participation.

A common classification distinguishes between implicit and explicit DSF. Implicit flexibility relies on indirect economic signals, such as time-of-use tariffs, critical peak pricing, or real-time pricing, encouraging consumers to adjust their behaviour voluntarily. Since the response is not contractually guaranteed, implicit flexibility can support load shaping and peak reduction, but is less suited to real-time system needs requiring firm activation [49]. Explicit flexibility, by contrast, involves a committed response from consumers or aggregators, usually associated with remuneration, penalties, prequalification, and measurement and verification procedures. It is therefore more suitable for balancing, congestion management, and other operational services where the delivered response must be dependable.

A related classification distinguishes between price-based and incentive-based DR schemes. Price-based schemes expose consumers to time-varying electricity prices and rely on their willingness and ability to modify consumption. Incentive-based schemes involve direct remuneration

for reducing, increasing, or modulating demand under predefined conditions, often through contracts, market bids, aggregators, or system-operator activation [50]. These approaches are complementary. Price-based schemes can improve demand elasticity and reduce peak demand, while incentive-based schemes are more closely aligned with dispatchable flexibility products.

DSF can also be classified according to the physical effect produced on the load profile. Curtailment reduces electricity consumption during a limited period, while load shifting moves consumption from one period to another, usually requiring later recovery of the deferred energy. Valley filling increases consumption during low-price or high-renewable-generation periods, whereas peak shaving reduces consumption during system stress or expensive periods. These actions differ in duration, recovery effects, user impact, and modelling requirements [48]. A complementary classification concerns service directionality. Depending on the adopted sign convention, upward flexibility may refer to actions that help address supply shortages, such as demand reduction or additional local injection, whereas downward flexibility may refer to actions that absorb excess supply, such as demand increase or reduced local injection. Because this terminology may vary between system-operator, aggregator, and end-user perspectives, the adopted convention should always be made explicit [47]. The main classification dimensions discussed above are summarised in Table 2.2, which highlights the operational interpretation of each category.

Table 2.2: Main classification dimensions of DSF.

Classification	**Main categories**	**Interpretation**
Activation mechanism	Implicit and explicit flexibility	Separates voluntary response to signals from committed and verifiable activation.
Economic signal	Price-based and incentive-based schemes	Distinguishes tariff- or market-price response from direct remuneration for specific actions.
Physical load effect	Curtailment, load shifting, valley filling, and peak shaving	Describes how the load profile changes and whether rebound or recovery effects may occur.
Directionality	Upward and downward flexibility	Identifies whether the action helps cover shortages or absorb excess supply, according to the adopted sign convention.
Operational attributes	Power magnitude, duration, response time, location, availability, controllability, and predictability	Determines whether flexible demand can become a reliable service for operation, aggregation, or markets.

Finally, DSF classification must account for the constraints of the resources that provide it. In residential buildings, flexible operation must balance grid-service provision with thermal comfort, occupant behaviour, and local technical limits [51]. More generally, flexibility may also be constrained by routines, production schedules, process continuity, service-quality requirements, or user acceptance. Therefore, classification is only the first step. The practical use of DSF in EMSs and aggregation frameworks requires context-specific modelling of availability, activation limits, rebound effects, baselines, and verification procedures.

2.2.2 Demand-Side Flexibility in B2B and B2C Contexts

DSF differs substantially across B2B and B2C consumers because electricity use is driven by different operational, economic, and behavioural factors. In industrial and commercial contexts, flexibility is often linked to production schedules, service continuity, equipment constraints, automation capabilities, and the economic value of modifying demand. In residential contexts, by contrast, flexibility is usually distributed across many small resources and is strongly influenced by comfort preferences, daily routines, user acceptance, and the willingness to delegate control. Therefore, B2B and B2C flexibility differ not only in scale, but also in predictability, controllability, aggregation needs, and the constraints to be respected [52].

In B2B industrial settings, the flexibility potential can be significant because individual facilities may contain large electricity-intensive processes. However, this potential often cannot be represented as a simple load reduction. Industrial flexibility must usually be embedded within production planning, maintenance schedules, safety constraints, quality requirements, and contractual obligations. Recent reviews of industrial DR in Europe emphasise that energy-aware scheduling, digitalisation, industrial aggregators, and sector-specific expertise are needed to translate technical flexibility into usable flexibility services [53]. Applied studies in power-intensive industries reinforce this point. In metal casting, for example, DSM requires optimisation models capable of coordinating multi-stage production processes while respecting holding times, material-flow constraints, furnace operation, and market requirements [54]. This illustrates why B2B industrial flexibility can be valuable, but also highly dependent on process-specific modelling.

Commercial B2B consumers present a different flexibility profile. Their loads are often less process-intensive than heavy industry, but their operation is still constrained by service quality, occupancy, customer activity, workplace requirements, and building comfort. HVAC systems are particularly relevant in commercial buildings because they can exploit building thermal mass to shift demand without directly interrupting business activity. Field demonstrations of MPC in large office buildings show that commercial HVAC flexibility can be activated in response to dynamic electricity prices while maintaining comfort constraints, although implementation may require significant data integration, customised control logic, and collaboration with facility managers [55]. As a result, commercial flexibility is technically promising, but its scalability depends on reducing engineering effort and adapting control strategies to heterogeneous building systems.

In B2C residential contexts, the main challenge is not the absence of flexible resources, but their small size, dispersion, and dependence on user behaviour. Heat pumps, EVs, smart appliances, thermal storage, and behind-the-meter batteries can provide meaningful flexibility when coordinated, but each individual household usually has limited capacity and uncertain availability. Field trials with residential air-source heat pumps show that third-party control can reduce demand during flexibility events, although successful deployment depends on automation, low user burden, comfort preservation, and the possibility of user overrides [56]. This is consistent with empirical evidence from residential heating flexibility, where households may be willing to participate, but acceptance can be constrained by interventions in the home, perceived loss of comfort,

and household-specific preferences [57]. Residential flexibility is therefore shaped not only by technical availability, but also by the social acceptance of automation.

These differences have direct implications for flexibility management and aggregation. Large industrial consumers may provide substantial flexibility at site level, but its mobilisation depends on sector-specific operational and market arrangements [53]. Smaller consumers, particularly at residential level, more often require aggregation to combine dispersed resources into a portfolio of sufficient scale and reliability [6]. However, aggregation does not remove the underlying constraints of each sector. Industrial portfolios still require process-aware scheduling, commercial portfolios require building-aware control, and residential portfolios require comfort-aware and behaviour-aware automation. Consequently, the practical value of B2B and B2C flexibility depends on how accurately availability, activation limits, rebound effects, and user constraints can be characterised.

2.2.3 Characterisation and Quantification of Demand-Side Flexibility

The effective use of DSF depends not only on identifying flexible resources, but also on describing their operational capabilities in a form suitable for control, aggregation, and verification. In practice, flexibility must be characterised by attributes such as power magnitude, response direction, activation time, duration, and possible rebound or recovery effects [6]. These characteristics determine whether a resource can support local optimisation, system operation, or market-based services.

From a practical perspective, flexibility is often expressed through modifications of the demand profile over time, including load shifting, peak shaving, and valley filling. A representative example is shown in Fig. 2.3, where the original demand profile is reshaped through a flexibility strategy. This representation links abstract flexibility concepts to observable load changes and highlights the intertemporal nature of many flexibility actions.

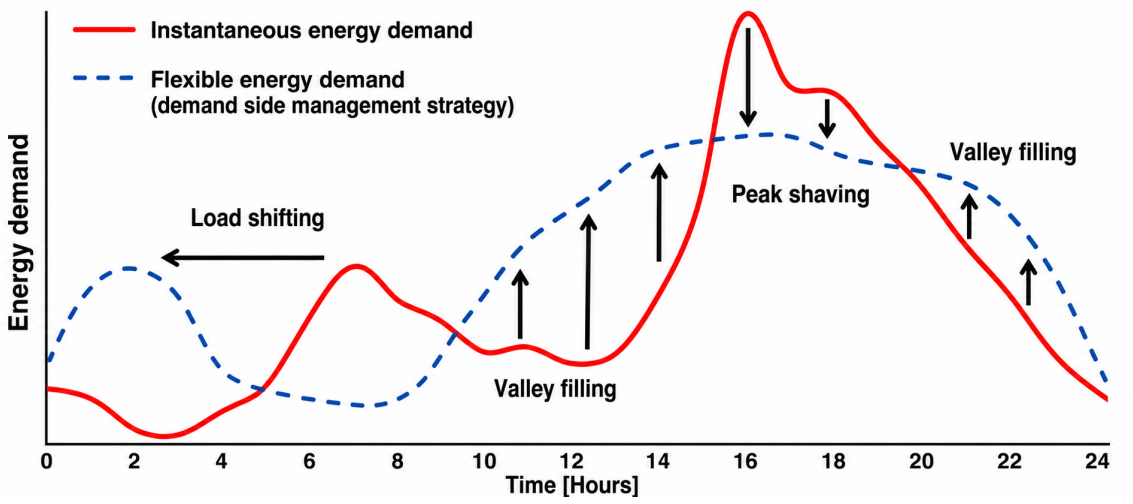


Figure 2.3: Demand profile reshaping through load shifting, peak shaving, and valley filling [58].

Characterising flexibility, however, requires more than visual inspection of reshaped load profiles. Some recent approaches represent flexibility through multidimensional indicators or flexibility envelopes that define the feasible operating region of a load, building, or portfolio over time [59]. These representations are useful because they capture the fact that an activation in one period may affect the flexibility available in subsequent periods, especially when thermal inertia, storage dynamics, or recovery requirements are present. In this sense, flexibility is not simply a static quantity, but a constrained temporal capability shaped by technical limits, user requirements, and system conditions. However, detailed flexibility envelopes can be costly to quantify and communicate, motivating compact data-driven representations for online applications [59].

A central challenge in flexibility quantification concerns the definition of the reference against which a flexibility action is measured. In many DR and market applications, the delivered flexibility is estimated as the deviation from a counterfactual baseline, representing the consumption that would have occurred without activation. Baseline construction may rely on historical averaging, day matching, weather matching, regression models, control groups, or more advanced data-driven approaches. Nevertheless, the baseline problem remains difficult because the counterfactual is never directly observed and may be strongly affected by occupancy, weather conditions, behavioural variability, and repeated flexibility events. Baseline quality therefore affects not only the estimated amount of delivered flexibility, but also the credibility of measurement and verification, since systematic over- or under-estimation can distort the apparent response and its economic valuation [60].

To address these limitations, recent studies have explored data-driven methods and dedicated flexibility indicators. In building applications, a key performance indicator (KPI) can be used to quantify flexibility according to the available data, the stakeholder perspective, and the intended use of the metric [61]. Some indicators require an explicit baseline, whereas others are designed to evaluate flexibility directly from measured operational data. More broadly, structured data-driven frameworks distinguish between retrospective analyses, which assess historical flexibility performance, and prospective analyses, which estimate future flexibility by combining measured data with predictive or physics-informed models [62]. This distinction is particularly relevant because the requirements for evaluating past flexibility are not the same as those needed to forecast future flexibility availability.

At portfolio level, flexibility cannot always be represented as the simple sum of individual capabilities, since different resources may exhibit distinct response times, recovery effects, uncertainty levels, and local constraints. In non-residential and multi-site applications, the literature combines measured data, building simulation, co-simulation environments, and optimisation-based methods to estimate site-level and aggregated flexibility more realistically [63]. Despite these advances, the lack of standardised flexibility metrics and high-quality datasets continues to limit comparison and validation across studies [61]. Baseline estimation and the measurement and verification of delivered flexibility remain additional challenges [60]. Consequently, the characterisation and quantification of DSF remain active research areas, particularly in contexts where flexibility must be measured, forecast, and aggregated.

2.2.4 Integration and Aggregation of Demand-Side Flexibility

The integration of DSF requires translating individual resource capabilities into a coordinated portfolio that can be managed as a reliable operational asset. This is particularly important for small consumers, whose limited and dispersed flexibility may prevent direct participation in markets or system services. Aggregation allows heterogeneous resources, such as HVAC systems, heat pumps, EVs, BESSs, and controllable appliances, to be pooled so that their combined response can provide higher capacity, smoother aggregate behaviour, and improved operational value [6]. However, aggregation does not eliminate local constraints: comfort limits, availability, rebound effects, and device-level restrictions must still be preserved when flexibility is activated.

Several modelling approaches have been proposed to represent aggregated flexibility in a compact and controllable form. Thermostatically controlled loads, for example, are often represented through virtual battery models, where the collective ability to increase or decrease demand is described analogously to the charging and discharging limits of a storage device [64]. Portfolio construction methods can also improve aggregation by selecting consumers through probabilistic forecast information, reducing forecast inaccuracy compared with random grouping at low aggregation levels [65].

Forecasting is therefore a central requirement for aggregated flexibility. Aggregators need to estimate both expected demand and the flexibility that can be reliably delivered under future conditions. This is challenging when new consumers have little historical response data, or when availability depends on time-varying variables such as weather, occupancy, user behaviour, equipment status, and previous activations. Online transfer learning has been proposed to address this issue by using knowledge from existing customers and updating the forecasting model as new information becomes available from recently added participants [66]. Such adaptive methods are relevant because aggregated flexibility must be updated as portfolio composition and operating conditions evolve.

The activation of aggregated flexibility also requires careful coordination before and after each event. If many resources are controlled simultaneously, rebound or recovery effects may create secondary demand peaks that partially offset the intended benefit. In virtual community studies, control-oriented aggregation frameworks have been used to coordinate residential buildings while considering comfort constraints and baseline-demand modification at portfolio level [67]. This illustrates that larger portfolios can increase controllability and capacity, but also require coordinated recovery and careful scheduling, while benefiting from participant diversity.

Finally, the practical value of aggregated flexibility depends on whether the portfolio can satisfy the requirements of the service being provided. Minimum bid sizes, activation durations, response times, symmetry requirements, and penalties for non-delivery can limit the ability of aggregators to convert local flexibility into dependable products [6]. Therefore, DSF aggregation should be understood as a coordination layer that combines forecasting, portfolio design, activation control, user constraints, verification, and risk mitigation, enabling flexible demand to be coordinated with energy-market participation, reserve provision, and procurement decisions.

2.3 Electricity Market Participation and Procurement Mechanisms

Electricity markets and procurement arrangements provide the economic and operational context in which distributed resources and portfolio-level energy-management decisions are valued. At European level, electricity trading is organised across forward, day-ahead, and intraday timeframes, while balancing markets and mechanisms are operated by TSOs closer to real time [68]. Day-ahead and intraday trading support scheduled electricity exchange before delivery, whereas balancing arrangements allow TSOs to correct imbalances and maintain the real-time balance between supply and demand [69].

For consumers, prosumers, and aggregators, this multi-timeframe structure creates different procurement options, revenue opportunities, and operational requirements. Day-ahead prices influence the scheduling of local generation, storage, EVs, and flexible demand. Balancing products, in turn, may remunerate both reserved capacity and the balancing energy delivered following activation. Long-term bilateral arrangements, particularly those established through a Power Purchase Agreement (PPA), can reduce exposure to short-term price volatility, provide greater price stability, and support renewable-energy procurement [70]. This section therefore reviews the mechanisms most relevant to the proposed aggregated EMS: day-ahead energy markets and price signals, balancing markets and reserve products with particular emphasis on mFRR, and bilateral energy procurement, including PPAs. The interaction between these mechanisms, the aggregator-managed EMS, and the underlying resource portfolio is illustrated in Fig. 2.4.

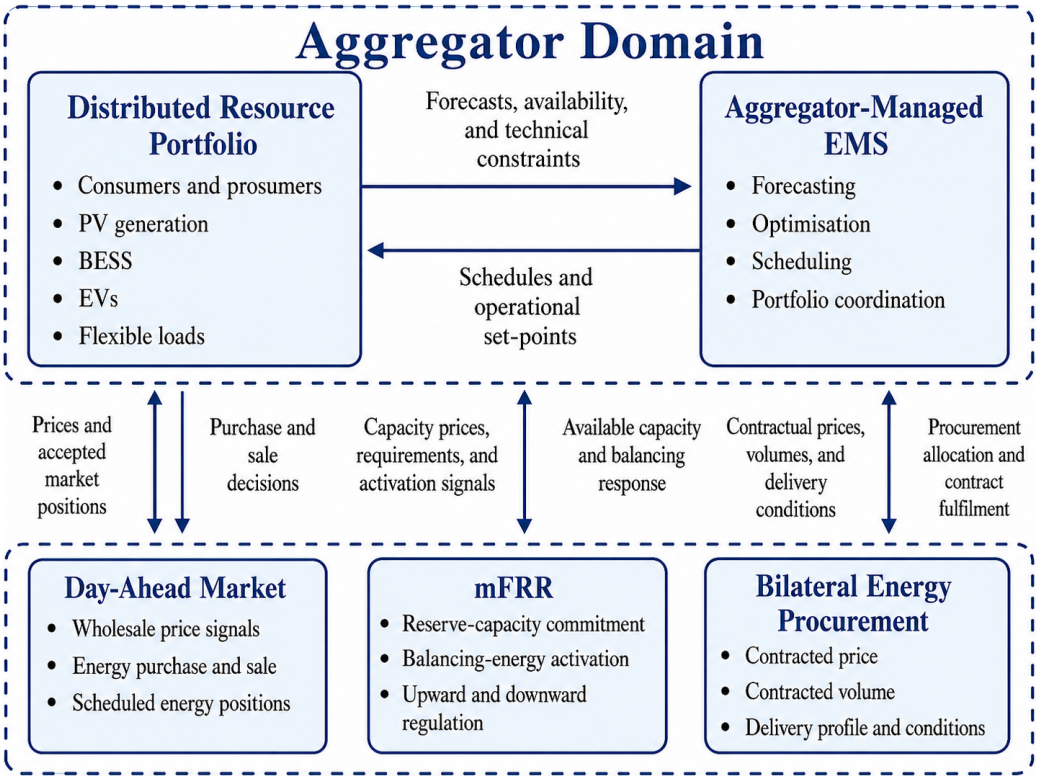


Figure 2.4: Aggregator-managed EMS interactions with resources and markets.

2.3.1 Day-Ahead Energy Markets and Price Signals

The DAM is one of the central short-term mechanisms for organising scheduled electricity exchange before physical delivery. In this market, participants submit purchase and sale bids for delivery on the following day, and the clearing process produces matched quantities, zonal marginal prices, and a day-ahead market schedule [71]. Because clearing occurs before delivery, the DAM provides an ex-ante wholesale price reference. It is, however, only one stage within the wider sequence of electricity market timeframes. Market participants may adjust their positions through intraday trading, while balancing arrangements address the remaining system imbalances closer to real time [68].

In Europe, day-ahead trading is coordinated through the Single Day-Ahead Coupling (SDAC), which aims to create a single pan-European cross-zonal day-ahead electricity market. The SDAC couples wholesale electricity markets across different bidding zones using a common algorithm, allocating scarce cross-border transmission capacity while accounting for transmission constraints and maximising social welfare [72]. This framework is supported by the European capacity allocation and congestion management rules, which provide the basis for harmonised day-ahead and intraday market coupling across bidding zones [73]. As a result, day-ahead prices reflect not only supply and demand bids within each bidding zone, but also the cross-border exchange capacity available within the coupled European market.

In the Iberian context, organised electricity trading takes place within the Iberian Electricity Market (MIBEL). Within this framework, Iberian Electricity Market Operator (OMIE) manages the day-ahead and intraday electricity markets, while Iberian Energy Derivatives Exchange (OMIP) manages forward-market activity [74]. OMIE publishes day-ahead prices for the Portuguese and Spanish bidding zones, together with the total power acquired in the Iberian day-ahead market and other market results [75]. These zonal prices constitute the main wholesale spot-market references for studies addressing energy procurement, surplus export, storage scheduling, and flexibility management in Portugal.

Although Portugal and Spain are integrated within the MIBEL, day-ahead prices should still be interpreted at the bidding-zone level. When the interconnection capacity between the two systems is sufficient, market coupling may lead to a common price in both zones. Conversely, when the interconnection becomes congested, market splitting occurs and the Portuguese and Spanish zones may clear at different prices. Energy Services Regulatory Authority (ERSE) reports that congestion episodes on the Portugal-Spain interconnection still occurred in 2024, despite the high degree of price convergence observed between both systems [76]. Under the European market-coupling framework, such price separation also generates congestion revenue through the cross-zonal capacity charge associated with the interconnected bidding zones [73], [76]. Therefore, for Portuguese consumers, prosumers, and aggregators, the Portuguese day-ahead price is the appropriate wholesale reference, rather than an Iberian average that may hide zonal price separation.

A recent development in European day-ahead market design is the transition towards shorter market time units. On 30 September 2025, for delivery from 1 October 2025, the European day-

ahead market moved from hourly to 15-minute trading intervals within the SDAC, allowing prices to be calculated every 15 minutes and expected generation and demand to be represented with finer temporal resolution [77]. This development is consistent with Regulation (EU) 2019/943, which requires trading intervals in the day-ahead and intraday markets to be at least as short as the applicable imbalance settlement period [78]. In the Iberian market, OMIE also implemented 15-minute products in the day-ahead market, allowing agents to submit purchase and sale bids in quarter-hourly blocks. Iberian intraday trading had already adopted this granularity in March 2025 [79]. Finer market time units allow variable renewable generation to be represented more accurately and enable flexible resources, such as BESSs, EVs, and controllable loads, to adjust their schedules at a finer resolution. This finer granularity also improves the alignment between market schedules and physical system conditions, helping to reduce imbalances and the associated costs of managing electricity flows [77]. Nevertheless, because the framework developed in this dissertation is evaluated at hourly resolution, the transition is treated as a recent market-design development rather than as a feature represented directly in the simulations.

For consumers and aggregators, the day-ahead price should not be confused with the final retail electricity price paid at the point of consumption. The day-ahead price is a wholesale reference, whereas the final consumer price depends on the supply contract, energy tariffs, network access tariffs, regulated charges, fees, taxes, and other commercial or regulatory components. In Portugal, electricity invoices include elements such as contracted power, energy prices, applicable tariffs, the unbundled value of the grid access tariff, fees, taxes, and costs of general economic interest [80]. For customers in the liberalised market, electricity prices are freely agreed between suppliers and their customers, while consumers remaining in the regulated market pay a transitional tariff set by ERSE. In both cases, all consumers pay network access tariffs established by ERSE [81].

The structure of these regulated and commercial components affects how price signals are perceived by final consumers. Network access tariffs include, among other elements, costs associated with the overall use of the electricity system, the transmission network, and the distribution networks [82]. Tariff options may also differentiate electricity prices according to time periods, such as peak, mid-peak, off-peak, and super off-peak periods, depending on the voltage level and tariff cycle applicable to each consumer [81]. Therefore, the wholesale day-ahead price is best understood as one component of a broader price environment, rather than as a direct representation of the final retail bill.

For the proposed EMS, this distinction has three implications. First, market-indexed procurement may follow the day-ahead wholesale price, while electricity supplied to final consumers may be priced through retail tariffs, supplier contracts, or regulated tariff structures. Second, surplus electricity may receive a specific export remuneration, while bilateral contracts, including PPAs, may introduce contractual prices and volume conditions that coexist with spot-market exposure. Third, the value of day-ahead price signals depends on how they interact with tariffs, export remuneration, contractual procurement, and local resource constraints. The day-ahead energy market considered in this dissertation values scheduled energy exchanges, whereas reserve availability and balancing-energy activation are treated through separate balancing arrangements.

2.3.2 Ancillary Service Markets and Reserve Products

Day-ahead and intraday markets establish scheduled electricity exchanges before delivery, but power-system operation also requires continuous real-time balancing between injections and withdrawals. In the European market sequence, balancing is treated as a separate market run by TSOs, allowing imbalances to be corrected when demand and supply do not match [68]. This physical need for balancing is reflected in system frequency, whose nominal value is set at 50 Hz in European synchronous areas [83]. When scheduled and actual injections or withdrawals diverge, balancing mechanisms and reserve products are therefore required to maintain frequency and operational security. Consequently, these mechanisms complement scheduled energy markets by providing the operational response that day-ahead and intraday trading cannot by themselves guarantee.

Within this framework, market-based balancing relies on distinct responsibilities and service-provision roles. A Balance Responsible Party (BRP) is responsible for the imbalances associated with its portfolio, whereas a Balancing Service Provider (BSP) provides balancing services to TSOs under the applicable terms and conditions for balancing [84]. The same European framework distinguishes between balancing capacity and balancing energy: balancing capacity refers to reserve capacity held available by a BSP, while balancing energy corresponds to the energy used by TSOs to perform balancing actions [84]. This distinction is important for flexible resources and aggregators, since participation in reserve products is not limited to real-time energy delivery. It may also require preserving available capacity before any activation occurs. The European balancing framework is therefore relevant for distributed flexibility, as it aims to enable new resources, including demand response and renewable generation, to participate in balancing markets [69].

Reserve products differ according to their activation speed, control logic, and operational purpose. Under the European system-operation framework, Frequency Containment Reserve (FCR) provides the initial containment layer for frequency deviations, while automatic Frequency Restoration Reserve (aFRR) and mFRR contribute to frequency restoration through automatic and manual activation, respectively. Replacement Reserve (RR) has historically provided a replacement function to restore reserve availability after previous activation [83]. European balancing integration has also been supported by common platforms for balancing-energy exchange: Platform for the International Coordination of Automated Frequency Restoration and Stable System Operation (PICASSO) for aFRR, Manually Activated Reserves Initiative (MARI) for mFRR, and Trans European Replacement Reserves Exchange (TERRE) for RR [85], [86], [87].

For EMSs and aggregators, this hierarchy means that the suitability of a flexible resource for a given reserve product depends not only on the available power, but also on its activation delay, ramping capability, availability requirements, and ability to sustain the required response [83]. These differences are illustrated in Fig. 2.5.

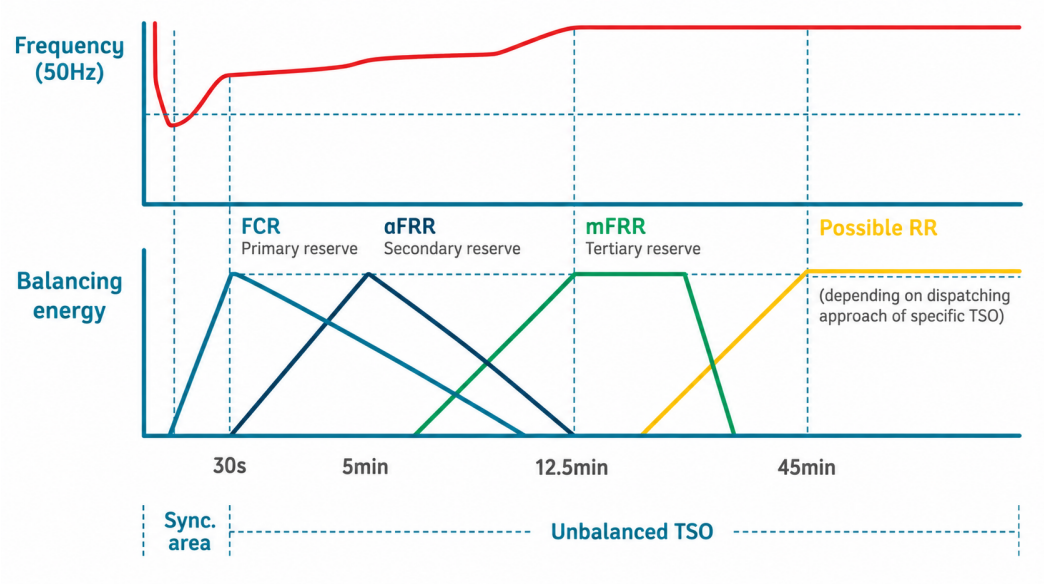


Figure 2.5: Temporal activation hierarchy of FCR, aFRR, mFRR and RR [88].

The hierarchy in Fig. 2.5 should be interpreted alongside recent changes in European balancing market integration. The TERRE project implemented the LIBRA platform for the exchange of RR balancing energy. Operations ended on 30 December 2025, after which the platform was decommissioned, closing the project at the end of March 2026 [87]. This transition was anticipated by the RR TSOs, which announced that it would no longer be feasible to exchange standard RR products under the new market-design timing requirements [89]. In Portugal, Redes Energéticas Nacionais (REN) reported that the last delivery period for trading RR energy on LIBRA occurred on 30 December 2025 and that, from that point onwards, it would no longer be possible to trade the RR product in the National Electricity System [90]. Therefore, while RR remains relevant for understanding the historical reserve hierarchy, its practical role in current Portuguese and European balancing-market participation has been reduced. This motivates a stronger focus on mFRR, which remains directly relevant for flexible portfolios because it combines reserve-capacity commitments with possible balancing-energy activation.

In Portugal, this focus on mFRR is reflected in the integration of REN into the MARI platform. On 27 November 2024, REN successfully accessed MARI, enabling the exchange of mFRR balancing energy through the European platform [91]. At national level, the *Manual de Procedimentos da Gestão Global do Sistema* distinguishes between two complementary dimensions of mFRR: the daily mFRR band, which corresponds to the capacity or availability component, and mFRR balancing energy, which corresponds to the energy activated for manual frequency restoration.

The daily mFRR band is addressed in Chapter XVI. Article 181 defines the scope of the service, including the methodology for determining band requirements, the contracting model, and the respective valuation [92]. The same article frames this band as a standardised capacity service aligned with the European methodology for standard balancing-capacity products approved

by European Union Agency for the Cooperation of Energy Regulators (ACER) through Decision No. 11/2020 [93]. Article 182 further establishes that the global system management function must ensure the contracting of the required daily mFRR band for the Portuguese load-frequency control area, and that each BSP with a contracted band must provide the corresponding capacity through mFRR energy offers [92]. The energy component is treated separately in Chapter XVII, where Article 194 defines the mFRR energy product, the methodology for determining mFRR needs, the mobilisation process, and the corresponding valuation [92]. The Portuguese framework therefore comprises two related but distinct mFRR components. The daily mFRR band represents contracted reserve capacity, which must be made available through corresponding direct-activation mFRR energy offers and may subsequently be activated as mFRR balancing energy [92].

The operational characteristics of these two components impose concrete requirements on flexible portfolios. For the daily mFRR band, Articles 183 and 184 state that participation requires market agents to be qualified as BSPs, and that the corresponding physical units must be enabled to provide mFRR energy [92]. For mFRR balancing energy, Article 196 requires market agents to obtain BSP qualification, while Article 197 requires the participating physical units to demonstrate the applicable technical and operational capabilities. The same article allows the aggregation of consumption, production, and/or storage installations below 1 MW, provided that the resulting aggregated physical unit has an offer capacity of at least 1 MW [92]. The market-based nature of the service is reinforced by the valuation rules: Article 191 values contracted daily mFRR band offers according to the price of the last accepted offer in the corresponding regulation direction, while Article 218 values activated mFRR energy using the corresponding marginal mFRR price according to the activation type, regulation direction, and delivery period [92]. The most relevant operational requirements are summarised in Table 2.3.

Table 2.3: Main operational requirements for Portuguese mFRR participation.

Operational aspect	**Daily mFRR band**	**mFRR balancing energy**
Service role	Capacity reserved before activation.	Energy activated for manual frequency restoration.
Minimum offer	1 MW minimum quantity; 1 MW resolution.	1 MW minimum quantity; 1 MW resolution.
Time granularity	Day-ahead procurement with 15-minute contracting and offer-validity periods.	15-minute market time unit; minimum delivery period of 5 minutes.
Activation	Capacity made available through direct-activation mFRR energy offers.	Manual scheduled or direct activation; total activation time of 12.5 minutes.
Direction and location	Upward or downward regulation in the Portuguese load-frequency control area.	Upward or downward activation in the Portuguese load-frequency control area.
Aggregation relevance	Helps portfolios meet minimum capacity and availability requirements.	Units below 1 MW may participate through aggregation if qualification requirements are met.
Valuation	Marginal price set by the last accepted offer in each contracting period and regulation direction.	Marginal pricing for each delivery period, differentiated between scheduled and direct activation.

These requirements show why aggregation is central to distributed mFRR participation. An EMS-managed portfolio can combine BESSs, EVs, flexible loads, and small generation units to meet minimum capacity, qualification, availability, and activation requirements that individual assets may not satisfy alone. At the same time, the separation between the daily mFRR band and mFRR balancing energy creates an operational trade-off, since capacity reserved for balancing services must remain available for possible activation instead of being freely allocated to energy arbitrage, self-consumption, electric-vehicle charging, or demand-side flexibility management.

2.3.3 Power Purchase Agreements and Bilateral Energy Procurement

PPAs are a prominent form of medium- or long-term bilateral energy procurement and provide a contractual complement to short-term electricity market participation. Whereas short-term market participation involves time-varying prices and operational commitments close to delivery, PPAs allow electricity producers and buyers to agree on contractual conditions over longer horizons. In renewable procurement, these contracts can provide greater revenue certainty for producers and price stability for buyers, reducing their exposure to short-term wholesale price volatility [70]. For portfolio-based energy management, their relevance extends beyond financial hedging, since contracted prices, volumes, and delivery profiles can influence the coordination of spot-market exposure, local generation, storage operation, and flexible demand. A simplified representation of the contractual, physical, and financial relationships involved in a physical renewable PPA is shown in Fig. 2.6.

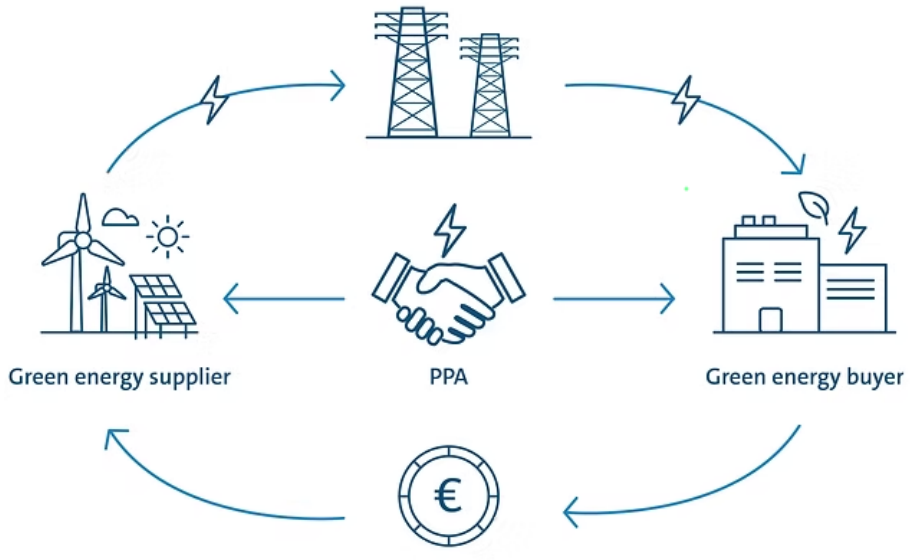


Figure 2.6: Simplified structure of a physical renewable PPA [94].

The structure of a PPA determines how price, volume, profile, balancing, and counterparty risks are allocated between the contracting parties. Physical PPAs involve the delivery of electricity, either directly on-site or through the public power system, and may require additional supply,

scheduling, market, or balancing arrangements. Virtual or financial PPAs, by contrast, are commonly structured as contracts for difference that settle the difference between an agreed strike price and a market reference price, while physical electricity procurement and imbalance management remain subject to separate arrangements. Beyond this distinction, recent reviews classify PPAs according to dimensions such as delivery location, delivery profile, pricing structure, buyer configuration, generation technology, contract duration, and risk allocation [95]. These characteristics are relevant because the same nominal contracted volume may imply substantially different operational and financial responsibilities for the buyer.

Renewable PPAs also introduce specific risks that must be considered when they are used as procurement instruments. In pay-as-produced contracts, the contracted volume follows the generation profile of the renewable asset, which may not coincide with the buyer's consumption profile. This creates volume and profile risk, since energy deficits may need to be covered in the spot market while surpluses may have to be resold under uncertain price conditions [96]. Moreover, the economic value of a renewable PPA depends on the relationship between the agreed contract price and the capture price of the contracted generation. The capture price represents the production-weighted average market value of the electricity generated over time [97]. Hedging the resulting joint price and production exposure is not straightforward, as empirical evidence shows that conventional futures-based strategies may provide only limited protection against the combined risks of renewable PPAs [98].

The policy relevance of PPAs has increased within the European electricity market design. Regulation (EU) 2024/1747 requires Member States to promote their uptake by removing unjustified barriers and disproportionate or discriminatory procedures or charges, thereby improving price predictability and supporting decarbonisation objectives [70]. ACER monitors the development of PPA markets through annual assessments, evaluates the need for voluntary contract templates, and addresses remaining barriers to market integration [99]. This direction was reinforced by Commission Recommendation (EU) 2026/917, which recommends measures to remove regulatory and non-regulatory barriers, facilitate voluntary market platforms and guarantee schemes, and improve access for smaller buyers through multi-buyer PPAs and demand aggregation [100].

Portugal has also established a dedicated framework for the registration and voluntary bilateral contracting of electricity. Decreto-Lei n.º 99/2024 establishes the legal basis for this activity and assigns ERSE responsibilities for its regulation, monitoring, and supervision [101]. Portaria n.º 367/2024/1 defines bilateral contracting as a medium- or long-term energy purchase agreement between a seller and a buyer, designates OMIP as the entity responsible for managing the activity, and requires it to operate through an electronic platform [102]. ERSE subsequently approved the *Manual de Procedimentos da Atividade de Registo e Contratação Bilateral de Energia Elétrica*, which establishes the terms and procedures applicable to the mandatory registration of bilateral contracts and to voluntary bilateral contracting through the platform [103]. This framework connects contractual procurement with regulatory supervision, market transparency, and the programming requirements of the Portuguese electricity system. It thereby places bilateral procurement within the broader operational coordination of the national electricity system.

From an EMS perspective, physical bilateral procurement arrangements, including physical PPAs, can be represented as contractual options defined by specific prices, volumes, durations, and delivery conditions. Virtual PPAs, by contrast, require a financial-settlement representation rather than a physical energy procurement channel. The operational value of physical contracts depends on their coordination with spot-market purchases, local generation, storage operation, flexible demand, and reserve provision. Bilateral procurement therefore adds a contractual layer that can reduce exposure to short-term price volatility while introducing profile mismatch, imbalance exposure, counterparty risk, and valuation uncertainty.

2.4 Aggregation Frameworks and Business Models

The previous sections examined the technical foundations of energy management, the characterisation of DSF, and the market and procurement mechanisms through which flexible resources can create value. However, the practical mobilisation of distributed flexibility also depends on the actors that connect consumers, prosumers, flexible assets, markets, and system operators. In this context, energy aggregators occupy a central position because they transform dispersed and heterogeneous resources into coordinated portfolios that can be represented, activated, verified, and commercialised on a larger scale [104].

Aggregation frameworks, therefore, add an organisational and commercial layer to the EMS and market mechanisms discussed earlier. They define who coordinates flexible resources, how responsibilities are allocated, how consumers interact with external parties, and how flexibility is converted into reliable services. Fig. 2.7 illustrates this coordination role by positioning the aggregator between distributed resources and the main institutional actors of the electricity system.

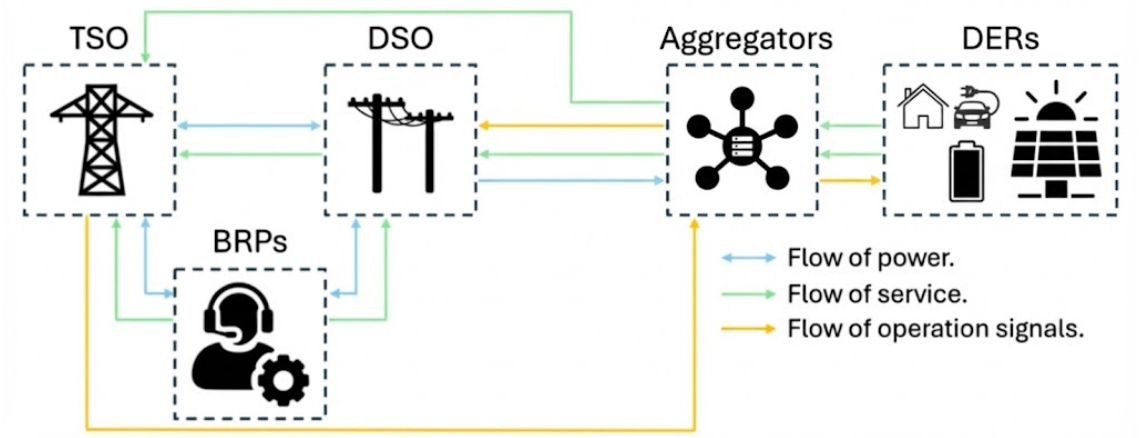


Figure 2.7: Coordination framework linking TSO, DSO, BRPs, aggregators, and DERs. Adapted from [105].

Fig. 2.7 distinguishes physical power flows, service exchanges, and operational signals between the main actors. While the DERs remain physically connected to the electricity network,

the aggregator coordinates their flexibility and represents the resulting portfolio capability towards system operators and other market actors.

This section reviews four complementary perspectives: aggregator roles, organisational structures, business models and revenue streams, and the digital platforms required to coordinate real-world flexibility portfolios.

2.4.1 Definition and Role of Energy Aggregators

In the European electricity framework, aggregation is formally understood as the function of combining multiple customer loads or generated electricity for sale, purchase, or auction in electricity markets. An independent aggregator is more specifically defined as a market participant engaged in aggregation that is not affiliated with the customer's supplier [106]. This legal framing is important because it treats aggregation not only as a technical optimisation problem, but also as a market activity associated with customer choice, third-party participation, and non-discriminatory access to flexibility resources. Recent European monitoring confirms that the legal recognition and practical eligibility of aggregators still vary across Member States, meaning that the same technical aggregation capability may face different market-access conditions depending on the national regulatory framework [107].

The need for aggregators is particularly clear for small, dispersed, or heterogeneous consumers. Individual residential prosumers, small commercial consumers, and many service-sector users may lack the volume, expertise, metering infrastructure, or administrative capacity required to participate directly in organised electricity markets or system-service mechanisms. Aggregators reduce these transaction barriers by pooling flexible loads, distributed generation, storage systems, EVs, and other DERs into portfolios that can satisfy minimum size, reliability, and coordination requirements [15]. This role is especially relevant for explicit DR involving small end-users, since the European regulatory framework creates opportunities for independent aggregators to combine the resources of multiple consumers and represent them in market-based mechanisms [108].

From an operational perspective, the aggregator performs a sequence of functions that go beyond the simple summation of individual schedules. These functions include customer enrolment, resource qualification, data collection, forecasting, and baseline construction. They also cover portfolio coordination, bid preparation, activation, monitoring, verification, settlement, and remuneration [31]. Their exact implementation depends on the regulatory context and on the flexibility product being delivered, but the underlying logic is common. The aggregator must convert uncertain and user-constrained local capabilities into a dependable portfolio-level response. Market commitments must therefore reflect the flexibility that underlying resources can reliably provide. This coordination requires technical and commercial integration while preserving consumer constraints [31].

A useful distinction can therefore be made between an EMS and an aggregator. An EMS is the technological and decision-making layer used to monitor assets, process forecasts, solve scheduling problems, and issue control actions. By contrast, the aggregator is the organisational and market-facing actor that uses such tools to coordinate resources, manage contracts, interact with

consumers, and communicate with external entities. In digitalised power systems, this distinction becomes particularly visible because aggregators may interface with local EMSs or HEMSs while also exchanging information with DSOs, TSOs, BRPs, retailers, and market platforms [109]. The EMS is therefore part of the aggregator's operational capability, but it does not replace the aggregator's broader commercial, regulatory, and coordination responsibilities.

These responsibilities create several practical and regulatory challenges. When an independent aggregator activates flexibility from consumers supplied by another market party, the resulting change in metered consumption can affect the balance position of the supplier or its BRP. This creates the need for mechanisms such as perimeter correction, baseline definition, and transparent data exchange to separate the aggregator's activation from the supplier's balance responsibility [16]. Financial compensation is another debated issue, since DR activation may alter the energy volume expected by the incumbent supplier. ACER has therefore emphasised that removing barriers to DR requires clearer rules for market access, energy allocation, compensation, baseline validation, and the participation of aggregated resources across wholesale and balancing markets [110].

Energy aggregators should therefore be understood as enabling actors that connect local flexibility with system- and market-level needs. Their value lies in organising participation, reducing transaction costs, managing information flows, and allocating responsibilities between consumers and external actors. Their practical role nevertheless depends on the organisational relationship between the aggregator, supplier, BRP, consumer, and system operators. These alternative arrangements are examined next.

2.4.2 Aggregation Paradigms and Organisational Models

Aggregation can be organised through different institutional and commercial arrangements, depending on who controls the flexible resources, maintains the relationship with the consumer, and assumes responsibility for market positions and imbalances. A first distinction can be made between integrated models, in which aggregation is performed by or in association with the consumer's supplier, and independent models, in which a third party manages flexibility separately from electricity supply. In integrated aggregation, supply, customer management, and flexibility contracting are combined within the same commercial arrangement, while balance responsibility may also be integrated depending on the applicable BRP structure. This simplifies coordination and avoids settlement with an external supplier, but concentrates energy supply and flexibility provision within the same actor [104]. Independent aggregation can strengthen competition and consumer choice, but requires explicit rules governing the interaction between the aggregator, supplier, and BRP.

Independent aggregation is not a single organisational form. In the contractual model, the aggregator and supplier rely on separate BRPs, with bilateral arrangements and information exchange used to address settlements and the imbalances caused by flexibility activation. In the delegated model, the supplier remains the sole BRP, while the aggregator commercialises flexibility at its own risk under predefined arrangements with the supplier. In the service-provider model, the

aggregator does not trade flexibility or assume balance responsibility, but provides another market actor with access to the coordinated resources [109]. These models determine who bears market and imbalance risk, who communicates with system operators, and how the consumer accesses the aggregation service.

The separation between electricity supply and flexibility management also creates an organisational problem around imbalance and compensation. When an independent aggregator activates DR, the metered consumption of the customer changes, even though the customer's supplier may have previously procured energy based on a different expected demand profile. Regulatory frameworks therefore need mechanisms to correct the supplier's balance perimeter and, where applicable, define compensation between the aggregator and the supplier [16]. Recent analysis of independent aggregation in the Nordic DAM confirms that the design of supplier compensation can affect welfare outcomes and remains a relevant market-design question [111]. For this reason, aggregator models cannot be evaluated only by their technical ability to pool resources. They must also be assessed according to how responsibilities, risks, and financial transfers are distributed among market actors.

VPPs and virtual aggregators represent another major aggregation paradigm. Their defining feature is the coordinated representation of distributed resources as a virtual portfolio that can interact with markets, system operators, or network constraints. The literature commonly distinguishes between commercial VPPs, which organise resources mainly for market participation and portfolio optimisation, and technical VPPs, which place stronger emphasis on network feasibility, ancillary services, and system-operation requirements [20]. This distinction is useful because VPPs are not only software platforms or optimisation models. They are also organisational structures that determine how distributed assets are grouped, controlled, and represented externally. At the same time, VPP operation increasingly involves multi-layer coordination between physical resources, communication systems, market interfaces, and organisational actors, which reinforces the need to treat VPPs as socio-technical aggregation structures rather than purely mathematical dispatch problems [19].

The internal governance of an aggregated portfolio may also be centralised, decentralised, distributed, or hybrid. Here, these terms refer specifically to the allocation of decision authority, information, and responsibility within the aggregation structure. Centralised arrangements provide stronger portfolio-wide visibility, whereas distributed arrangements preserve more local autonomy and privacy but require additional coordination [20]. The broader technical implications of these control structures were discussed previously in the review of EMS architectures.

Energy communities introduce a more local and collective form of aggregation. Renewable Energy Communities and Citizen Energy Communities are associated with shared ownership, local value creation, self-consumption, and collective participation, rather than with a purely commercial intermediary model of aggregation [112]. In this context, an aggregator or community manager may coordinate shared resources and optimise collective objectives, but participants may still retain local preferences and individual economic interests. Recent work on renewable energy communities illustrates this tension through structures where a central aggregator coordinates the

community while local participants may behave cooperatively or non-cooperatively depending on their objectives [32]. Community aggregation therefore differs from conventional portfolio aggregation because its organisational logic includes local governance, member participation, and social or environmental objectives.

The organisational paradigms discussed above, which may overlap in practice, are summarised in Table 2.4, highlighting their coordination logic and the principal trade-off associated with each model.

Table 2.4: Main organisational paradigms for energy aggregation.

Paradigm	Coordination logic	Main organisational trade-off
Integrated or supplier-led aggregation	The supplier or retailer coordinates flexibility and manages the consumer relationship, while balance responsibility may also be managed within the same commercial arrangement.	Simpler settlement and clearer responsibility allocation, but weaker separation between energy supply and flexibility services.
Independent aggregation	A third-party aggregator manages flexibility separately from the consumer's electricity supplier.	Greater competition and consumer choice, but additional requirements for imbalance correction, compensation rules, and data exchange.
Service-provider aggregation	The aggregator provides flexibility coordination or activation services to another market actor without full market exposure.	Lower wholesale-market risk, but narrower role in portfolio commercialisation and value capture.
VPP or virtual aggregator	Distributed resources are represented as a coordinated virtual portfolio for market, system-service, or network-facing operation.	Portfolio-wide coordination and access to multiple services, but more complex control authority, communication, and resource coordination.
Energy community aggregation	Local participants coordinate generation, consumption, storage, and flexibility through a community-oriented structure.	Potential for local value creation and consumer participation, but more complex governance and incentive alignment.

Among the paradigms summarised in Table 2.4, the proposed framework is operationally closest to a centrally coordinated VPP or virtual aggregator, with individually represented consumer resources coordinated at portfolio level for market and system-service participation. Community governance, shared ownership, and internal energy sharing are outside the model scope.

Overall, aggregation arrangements should be designed according to portfolio size, consumer type, regulatory responsibilities, data availability, control authority, and the degree of autonomy granted to participants. Larger B2B consumers may support more direct, contractual, or technically customised aggregation arrangements, whereas B2C consumers generally require models that reduce complexity and delegate market interaction to suppliers, independent aggregators, service providers, or community structures. The choice of organisational model therefore shapes how flexibility can be accessed and coordinated. Its economic viability depends on the business models, revenue streams, and value-allocation mechanisms discussed in the following subsection.

2.4.3 Business Models, Revenue Streams and Market Coordination

The organisational models discussed in the previous subsection define how aggregation is structured, but the viability of an aggregator ultimately depends on how distributed flexibility is

converted into economic value. Aggregator business models are shaped by two complementary dimensions: the value proposition offered to consumers and the external opportunities available to the aggregated portfolio. In residential and service-sector contexts, aggregators may reduce energy costs, simplify access to complex markets, manage flexible assets on behalf of end-users, or provide service-based contracts that reduce the need for direct consumer involvement [30]. Energy-as-a-Service approaches illustrate this evolution by replacing conventional volumetric retail logic with models in which prosumers pay for representation and optimisation in wholesale markets, giving aggregators a stronger operational and market-coordination role [113].

Commercial viability can be strengthened by combining multiple sources of value rather than relying on a single revenue stream. Energy cost optimisation, reserve provision, DR incentives, local flexibility procurement, congestion management, and internal portfolio balancing can complement each other when they are coordinated within the same portfolio. This value-stacking logic is particularly important because the availability of distributed flexibility is often variable, user-constrained, and uncertain at the individual-consumer level. By combining flexible demand, storage, EVs, PV generation, and other distributed resources, aggregators can increase the reliability of portfolio response and improve their ability to offer flexibility across different operational contexts [30].

Multi-market coordination is one of the clearest examples of this value-stacking logic. A prosumer aggregator can coordinate flexible consumption, PV generation, EVs, and thermostatically controlled loads to prepare bids that combine energy scheduling with secondary reserve availability [114]. In this setting, the economic role of aggregation is not limited to shifting consumption away from expensive periods. It also involves deciding how much flexibility should be reserved for system services, how much should be used for energy-cost reduction, and how uncertainty in consumer behaviour or renewable generation should be managed. Aggregator-managed portfolios may also create value internally, for example through local exchanges or community-level transactions before interacting with external markets. Holistic formulations of the energy aggregator problem show that internal trading, local market interaction, and wholesale trading can be jointly represented, although the resulting value depends strongly on the assumed price structures, trading fees, and access to each market layer [33].

The consumer-facing side of the business model is equally important because flexibility arrangements must be attractive and understandable to consumers while providing a sufficiently reliable response to the aggregator. Simple rewards or tariff discounts can lower the participation barrier for end-users, but they may not provide enough incentive for dependable response during high-value periods. More advanced load-aggregator models therefore coordinate external bidding with internal incentive-pricing decisions. Post-clearing pricing adjustments and the selection of users with more reliable responses can help manage uncertainty and reduce non-delivery risk [115]. Value distribution is therefore not merely a commercial detail. It directly affects the quality, predictability, and risk profile of the flexibility product offered by the aggregator.

Aggregator revenues also depend on the coordination framework through which flexibility is procured and activated by system operators or market actors. System-service procurement mecha-

nisms require rules for bid submission, activation priority, remuneration, and settlement, especially when the same resource could potentially serve both local network needs and system-level balancing requirements. The OneNet market-design framework shows that different TSO-DSO coordination architectures can be used to organise system-service procurement, including arrangements in which local and central markets interact through bid-forwarding mechanisms or priority rules [12]. For aggregators, these designs influence not only where revenue can be obtained, but also which operational constraints, activation risks, and settlement responsibilities accompany each revenue opportunity.

Despite these opportunities, the commercial scalability of aggregator business models remains uncertain. ACER highlights that aggregated demand-side resources still face barriers related to market access, baseline validation, compensation rules, data exchange, and settlement procedures [110]. Recent reviews of flexibility markets also argue that many studies focus on market design or optimisation results while giving less attention to the full cost stack required for implementation, including metering, communication infrastructure, information and communication technology systems, customer acquisition, transaction costs, and operational management [116]. This limitation is important because an economically attractive scheduling strategy may not translate into a viable business model once implementation costs, regulatory uncertainty, and consumer engagement requirements are considered. Operationalising these business models at scale also requires suitable digital platforms and implementation conditions, which are discussed in the following subsection.

2.4.4 Digital Platforms and Real-World Implementations

The previous subsection showed that aggregation business models depend on the ability to coordinate several sources of value. In practice, this coordination cannot be achieved only through optimisation models or contractual arrangements. It also requires digital infrastructures capable of connecting assets, consumers, aggregators, system operators, and market platforms. Digital flexibility platforms therefore support the operational chain through which distributed resources are enabled, qualified, aggregated, activated, measured, settled, and remunerated. This has been described as a flexibility-centric value chain, in which digital platforms coordinate the activities and information flows required to transform local flexibility into services that can be used by markets, system operators, or energy communities [117]. From an implementation perspective, these platforms can act as administrative coordinators, market intermediaries, or marketplace operators, depending on whether they mainly organise data exchange, connect buyers and sellers, or directly support trading and settlement [118].

Interoperability is one of the main requirements for scaling these platforms. Aggregators typically need to exchange data with local controllers, smart meters, market platforms, DSOs, TSOs, and sometimes retailers or BRPs. If each platform relies on proprietary data formats or closed interfaces, the cost of connecting new assets and new market channels increases significantly. Local flexibility market design studies therefore emphasise the role of multi-layer architectures, such as those based on the Smart Grid Architecture Model, to structure business, functional, information,

communication, and component layers in a coherent way [119]. At the technical level, open communication protocols, standardised data models, and secure application programming interfaces are important to reduce vendor lock-in and allow flexibility providers to connect heterogeneous devices and portfolios to different market environments [118].

Digital implementation also depends on the standardisation of the flexibility products being traded. From the platform perspective, bids must be described through attributes such as location, direction, capacity, duration, activation time, availability windows, and verification rules. The OneNet framework highlights that standardised active and reactive power products can reduce bidding complexity and support interoperability across the TSO-DSO-consumer value chain [120]. A common product definition also provides a consistent basis for comparing offers from heterogeneous resources and matching their capabilities with the specific needs of system operators. This is particularly relevant when distributed flexibility can be used to address both local network constraints and system-level services. In such cases, digital platforms must support coordination rules that avoid conflicting activations, double procurement, or the forwarding of flexibility bids that are not feasible from a distribution-network perspective. Recent DSO-oriented practice reports also point to the growing importance of flexibility information systems and controllable-unit registries for managing prequalification, activation, data exchange, and settlement at scale [121].

Measurement and verification are central to trust in digital aggregation. Where flexibility delivery is assessed against a baseline, platforms must compare the activated response with the expected consumption or generation pattern to quantify the delivered response. This requires smart meter telemetry, baseline methodologies, validation routines, and settlement procedures that are transparent enough to be accepted by consumers, aggregators, and system operators. Local flexibility market implementations, therefore, include measurement and verification as a core platform function rather than as a secondary administrative step [119]. Inaccurate baselines or insufficient telemetry can lead to incorrect remuneration, unfair penalties, or disputes between flexibility providers and buyers, which makes settlement design a key part of platform credibility [121].

Several real-world platforms illustrate how these concepts are already being implemented. Piclo operates as an independent marketplace for energy flexibility, connecting asset owners and flexibility providers with buyers across different opportunities [122]. NODES provides a trading platform through which flexibility service providers can sell services to system operators, including reservation and activation of flexibility [123]. GOPACS provides another example, as a joint initiative by Dutch grid operators focused on coordinating market-based solutions for congestion management [124]. In Portugal, the FIRMe project, launched by E-REDES, is particularly relevant because it tests a local flexibility market regime and seeks to onboard flexibility service providers in a national distribution grid context [125]. These examples show that digital platforms are not uniform: some act as independent marketplaces, others as grid-operator coordination tools, and others as national or regional pilots designed to test procurement, qualification, and user engagement.

Despite this progress, real-world implementation still faces significant barriers. Cybersecurity and data privacy become increasingly important as platforms rely on cloud-based services,

application programming interfaces, smart-meter data, and automated dispatch. Scalability also depends on standardised procedures, consumer onboarding, device compatibility, reliable communication, and clear responsibilities between flexibility providers and system operators [118]. Local flexibility markets remain fragmented across platform designs, data models, and coordination procedures [119]. The practical challenge is therefore to embed profitable aggregation strategies in digital infrastructures that are interoperable, secure, transparent, and acceptable to market actors and end-users.

2.5 Comparative Review and Research Positioning

Following the review of EMSs, DSF, electricity markets, bilateral procurement, and aggregation structures, this section examines how these elements have been combined in optimisation-based operational studies. The objective is not to provide an exhaustive mapping of the literature, but to compare a focused set of contributions that are directly relevant to the modelling scope and research positioning of this dissertation.

The comparison covers aggregator- and VPP-based operation, energy communities, and flexible portfolios of consumers or prosumers. Attention is given to the represented resources, market and procurement mechanisms, treatment of uncertainty, temporal coordination, and preservation of individual technical constraints. This comparison identifies the modelling dimensions that are already well established and those that remain only partially integrated within a single aggregated EMS framework.

2.5.1 Scope and Selection Criteria

The reviewed papers were identified through a targeted search of peer-reviewed literature using IEEE Xplore, ScienceDirect, SpringerLink, Scopus, Web of Science, and Google Scholar. The search considered combinations of terms related to aggregators, VPPs, energy communities, DSF, storage, electric mobility, energy and reserve markets, rolling operation, and bilateral or PPA-based procurement. The final sample comprises ten studies selected for their relevance to the operational decisions represented by the MILP-based aggregated EMS developed in this dissertation.

Selection favoured operational optimisation studies that represent coordinated consumers, prosumers, or distributed resources and address at least one relevant modelling dimension. These dimensions include multi-resource scheduling, energy or reserve market participation, demand flexibility, individual or network feasibility, uncertainty, real-time redispatch, and contractual procurement. Conceptual, regulatory, and review-based publications were not included in the comparative sample.

The review was not restricted to deterministic formulations. Stochastic, robust, bi-level, decentralised, and rolling-horizon approaches were retained when they provided relevant comparisons for specific modelling layers. The selected studies therefore combine close comparators

with complementary references addressing flexibility-band construction, reserve operation, uncertainty, real-time redispatch, and PPA-based scheduling.

2.5.2 Scientific Papers Review

The selected studies address different aspects of aggregator, VPP, energy-community, and flexible-portfolio operation. Each contribution is examined in terms of its optimisation approach, represented resources, market scope, temporal structure, and main differences from the operational framework developed in this dissertation.

A deterministic MILP model for the day-ahead operation of an energy aggregator is developed in [33]. The formulation manages individually represented prosumer households equipped with PV, household batteries, EVs, time-shiftable loads, inflexible demand, and detailed power-to-heat systems. Energy can be traded through wholesale and local markets or internally between households, with the objective of maximising their combined trading profit. The model provides a detailed representation of the interactions between heterogeneous resources and trading channels, while external flexibility and reserve services remain outside its scope.

A community-centred formulation is presented in [32], where a centralised MILP-based EMS manages a renewable energy community with residential, commercial, and industrial users. The community includes renewable generation, BESSs, EVs, and vehicle-to-building operation. The centralised EMS maximises shared energy while reducing BESS and EV battery cycling, and communicates reference grid-exchange profiles to the local EMSs. The analysis also assesses the effect of cooperative and non-cooperative user behaviour. Its scope is therefore centred on user-level coordination and energy sharing rather than external energy-market bidding or reserve provision.

Decentralised flexibility aggregation is addressed in [126]. In the proposed MILP framework, each unit determines a preferred operating profile and positive and negative reserve margins that remain feasible under its technical and user constraints. The approach is demonstrated using shiftable and thermostatically controlled loads, EVs, storage systems, PV, and non-controllable devices. The aggregator receives only the resulting profiles and reserve margins, which can be combined without access to each unit's internal model. This preserves local feasibility and information privacy, but does not centrally co-optimise the underlying resources and market layers.

Joint participation in day-ahead energy and secondary reserve markets is examined in [114]. The two-stage stochastic formulation coordinates PV, EVs with Vehicle-to-Grid (V2G) capability, thermostatically controlled loads, and inflexible demand. It jointly determines energy and reserve-band bids and accounts for band utilisation, energy imbalances, and reserve non-delivery across day-ahead and real-time stages. Uncertainty in generation, consumption, outdoor temperature, occupancy, and user preferences is represented through scenarios. The model offers a close methodological comparison for the joint valuation of energy and reserve availability, although the portfolio does not include dedicated stationary BESSs and is not updated through rolling re-optimisation.

Manual reserve provision by small prosumers is considered in [127]. The portfolio combines PV, V2G-enabled EVs, thermostatically controlled loads, shiftable loads, and inflexible consumption. A two-stage stochastic formulation defines day-ahead energy and tertiary reserve bids while accounting for real-time activation, energy imbalances, and reserve non-delivery. Unlike the secondary reserve formulation, the tertiary reserve product is remunerated through activated reserve energy rather than through a separate capacity-band payment. Shiftable loads provide an explicit representation of demand-side scheduling, while stationary BESSs and contractual procurement are not included.

Distribution-network feasibility is incorporated in [128]. The proposed decentralised strategy uses the alternating direction method of multipliers to decompose the problem into aggregator subproblems and a DSO subproblem based on alternating-current optimal power flow. The aggregators coordinate EVs, PV, and inflexible demand, while the DSO evaluates the feasibility of the resulting energy, upward-reserve, and downward-reserve operating scenarios. This structure preserves the independent roles and private information of the aggregators and the DSO. The contribution is focused on the formation of network-feasible bids, with real-time delivery and redispatch remaining outside the scope of the study.

A bi-level stochastic-robust framework for a VPP incorporating local energy communities is proposed in [129]. At the upper level, the VPP maximises its expected profit through energy and locational-flexibility trading. At the lower level, community energy managers minimise the costs of participating resources in response to dynamic retail prices. The portfolio includes EVs, thermostatically controlled loads, PV, microturbines, and inflexible demand. The resulting problem is reformulated as a single-level MILP and represents day-ahead bidding in the energy and flexibility markets, upward and downward flexibility, energy imbalances, and uncertainty in flexibility delivery. The flexibility product is market- and location-oriented rather than a frequency-reserve product, and coordination is induced through pricing instead of direct central dispatch.

The influence of storage on demand-response aggregation is analysed in [130]. A robust optimisation model maximises the profit of a demand-response aggregator operating in day-ahead and balancing markets under uncertain prices. Residential, commercial, and industrial users are represented together with rooftop PV, while demand response is mobilised through time-of-use and reward-based programmes. An aggregator-owned storage unit supports energy trading and reduces exposure to adverse market conditions. The storage system is not distributed among individual consumers, and demand response is represented through price and incentive programmes rather than explicit energy-neutral load shifting. Balancing-market transactions are also modelled without a separate reserve-capacity product.

Storage-based reserve scheduling and redispatch are examined in [131]. The proposed framework comprises two complementary models for a hybrid wind-battery VPP participating in day-ahead energy and secondary regulation reserve markets. A two-stage stochastic MILP determines the energy and reserve schedule using an hourly resolution, a look-ahead period beyond the market horizon, and forecasts of real-time reserve utilisation. Energy imbalances and battery degradation as a function of discharge depth and rate are also represented. A second deterministic MILP per-

forms hourly rolling redispatch using the schedule and storage opportunity cost obtained from the first model. This framework provides the closest comparison for look-ahead and rolling redispatch, but concerns a single hybrid generation-storage asset rather than a portfolio of consumers.

Finally, contractual and merchant market participation are combined in [132]. The study considers a semi-contracted and semi-merchant PV-BESS asset operating under a Pay-as-Delivered corporate PPA. A probabilistic neural network defines a secure contractual delivery profile, while a MILP determines BESS sizing, scheduling, and energy allocation among the PPA, the day-ahead market, the upward integrated scheduling process, and the upward mFRR segment. The formulation evaluates project value through contractual and market revenues while accounting for investment, operating, and non-delivery costs. It demonstrates how storage can be coordinated between contractual commitments and merchant market exposure, although the perspective is that of a renewable producer selling through a PPA, rather than an aggregator procuring energy for a consumer portfolio.

Taken together, the reviewed studies cover detailed multi-resource scheduling, community energy sharing, decentralised flexibility aggregation, energy and reserve bidding, network feasibility, uncertainty-aware demand response, rolling redispatch, and contractual market participation. These dimensions are nevertheless distributed across different portfolio structures and optimisation frameworks. The following subsection compares them systematically and identifies which combinations remain only partially represented in the literature.

2.5.3 Overview of Reviewed Contributions and Research Positioning

The selected studies address complementary operational dimensions of aggregated energy management, but combine them to different extents. Table 2.5 compares the dimensions that are most directly related to the integration scope of the proposed EMS. The comparison supports the research positioning of this dissertation without seeking to rank the selected studies.

Table 2.5: Comparison of selected operational optimisation studies.

Ref.	MILP	Port.	BESS	EV	DSF	DAM	Res.	Act.	Bil.	Roll.	Het.	Ind.
[32]	✓	✓	✓	✓	-	-	-	-	-	-	✓	✓
[33]	✓	✓	✓	✓	✓	✓	-	-	-	-	✓	✓
[114]	✓	✓	-	✓	✓	✓	✓	✓	-	-	✓	✓
[126]	✓	✓	✓	✓	✓	○	✓	✓	-	-	✓	✓
[127]	✓	✓	-	✓	✓	✓	✓	✓	-	-	✓	✓
[128]	○	✓	-	✓	-	✓	✓	○	-	-	✓	✓
[129]	✓	✓	-	✓	✓	✓	✓	✓	-	-	✓	✓
[130]	✓	✓	✓	-	✓	✓	-	-	-	-	✓	○
[131]	✓	-	✓	-	-	✓	✓	✓	-	✓	-	-
[132]	✓	-	✓	-	-	✓	✓	○	✓	-	-	-
This dissertation	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: ✓ = addressed; ○ = partially addressed; - = not addressed.

Table 2.5 uses the following abbreviations: Port., portfolio coordination; Res., external reserve or flexibility provision; Act., physical activation, redispatch, or activation-related economic valuation; Bil., bilateral or contractual decisions; Roll., rolling-horizon re-optimisation or redispatch; Het., participant heterogeneity; and Ind., individual technical feasibility. DSF excludes flexibility provided exclusively by EVs. A check mark denotes that a dimension is addressed, although its scope and detail may differ across studies. For this dissertation, Act. refers to expected activation revenues rather than physical activation.

The comparison shows that the selected literature covers most of the operational dimensions relevant to aggregated energy management, but rarely combines them within the same framework. Community- and portfolio-oriented formulations provide detailed representations of distributed resources and preserve user- or asset-level constraints. Their scope is nevertheless centred mainly on energy sharing, energy trading, internal coordination, or decentralised flexibility aggregation, rather than on the joint treatment of reserve provision, activation-related effects, rolling operation, and contractual procurement.

Market-oriented aggregator models provide a complementary perspective by combining day-ahead energy decisions with reserve or flexibility provision and the economic consequences of activation or imbalance settlement. These formulations represent flexible prosumer portfolios in considerable detail, but generally exclude dedicated stationary BESSs, bilateral procurement, or rolling re-optimisation within the same consumer-level framework.

Other dimensions are addressed mainly in more specialised system configurations. Rolling redispatch and activation-aware storage scheduling are developed for hybrid wind and battery VPPs, whereas bilateral and PPA-based operation is examined for utility-scale renewable energy and BESS assets. These studies provide close comparisons for particular modelling layers, but their system scope differs from that of an aggregated consumer portfolio with individual demand, storage, electric-mobility, and flexibility constraints.

The main research gap emerging from the comparison is therefore not the absence of the individual modelling components, but their limited integration within a single consumer-portfolio EMS. Existing studies tend to address multi-resource coordination, reserve participation, activation-related effects, rolling operation, individual feasibility, or contractual procurement separately, or in system contexts that differ from aggregated demand-side portfolios.

In this context, the proposed work is positioned as an integrated deterministic MILP-based EMS for aggregated consumer portfolios. It preserves individual consumer-level feasibility while coordinating PV generation, individual BESSs, EVs, explicit DSF, DAM participation, portfolio-level mFRR provision, expected activation revenues, and bilateral procurement within a rolling-horizon framework. By integrating these dimensions within the same operational framework, the proposed EMS addresses a modelling space that remains only partially integrated across the reviewed literature.

Chapter 3

Aggregated Energy Management System Model

This chapter presents the aggregated EMS model developed and extended in this dissertation for the coordinated operation of heterogeneous consumer portfolios. The framework solves a deterministic MILP problem over a rolling horizon, preserving the individual physical constraints of each consumer while coordinating selected market-oriented decisions at portfolio level.

The chapter begins with an overview of the framework structure and aggregation logic, followed by the description of the main resource models, covering BESS operation, data-driven DSF, and EV scheduling, with the EV model additionally incorporating rolling-horizon continuity for overnight sessions. These components are then integrated into the complete optimisation model, whose objective function and constraint blocks are presented by component. The chapter concludes with the computational implementation and result-processing workflow.

3.1 EMS Framework

The EMS considered in this dissertation is a modular optimisation framework designed for the coordinated operation of electricity consumer portfolios. It supports heterogeneous portfolio compositions, accommodating consumers with different demand profiles, asset configurations, tariff structures, flexibility availability, and market participation potential, while preserving the individual characteristics of each participant within a unified modelling environment.

The EMS builds upon a pre-existing optimisation tool developed at INESC TEC, originally applied to single-site energy and reserve scheduling [133], and extends it with additional modelling and operational layers. At its core, the model solves a deterministic MILP problem over a rolling time horizon. This formulation captures the main intertemporal and operational constraints of the managed resources while remaining computationally tractable for repeated portfolio-level scheduling. Dedicated modules translate input profiles, asset parameters, market signals, flexibility estimates, EV data, and bilateral contract limits into the parameters, variables, and constraints used by the optimisation model.

Fig. 3.1 illustrates the general structure of the proposed EMS across four layers: input data, portfolio resources, aggregated EMS, and outputs.

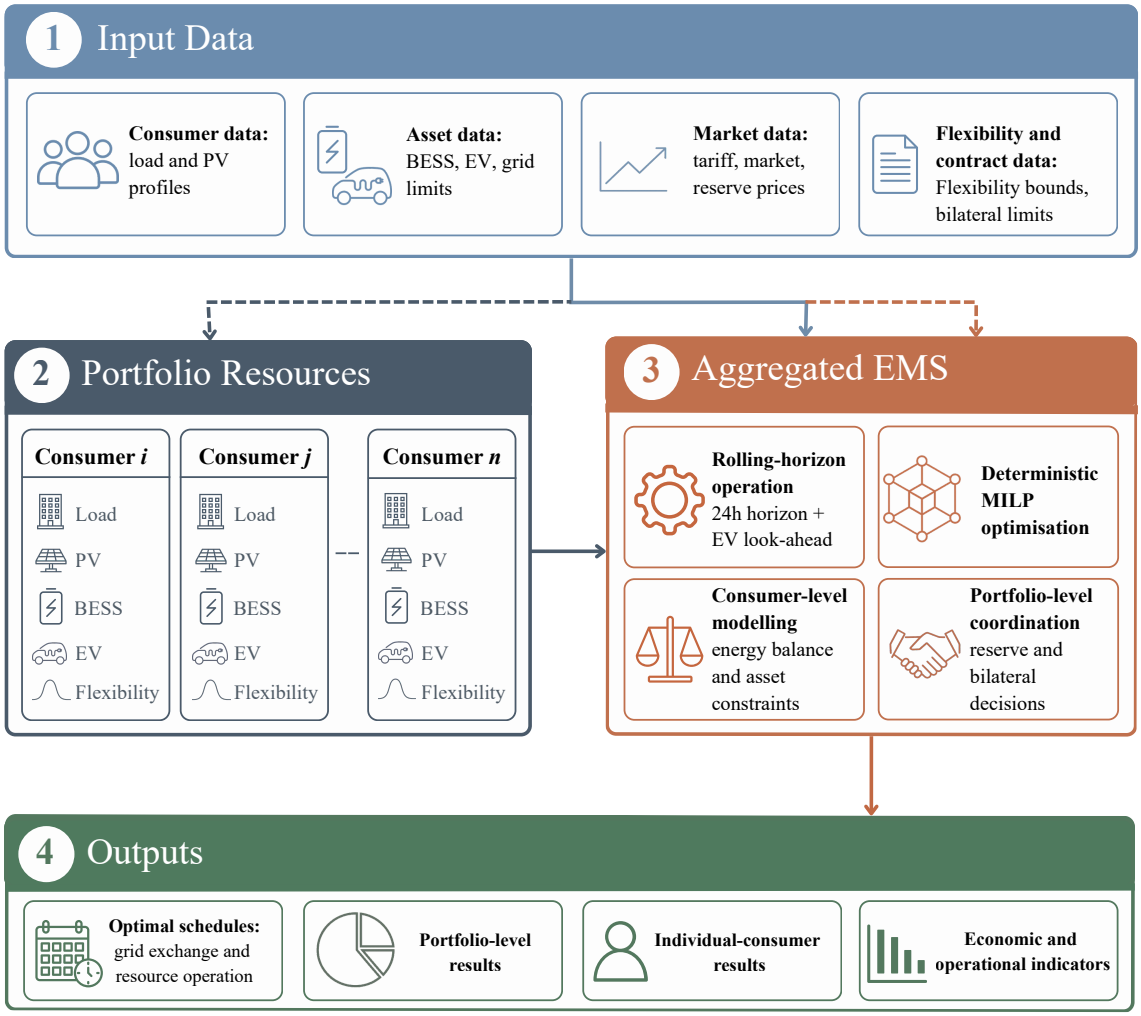


Figure 3.1: General structure of the proposed EMS framework.

A defining feature of the proposed EMS is that individual feasibility is preserved while selected decisions are coordinated at portfolio level. Each consumer is represented through its own load, PV generation, local assets, grid connection limits, EV parameters, and flexibility bounds, without reducing the portfolio to a single equivalent consumer. Market-oriented decisions such as reserve provision and bilateral procurement are managed collectively by the aggregator.

The modular design also supports a progressive evaluation of the available functionalities. Resource models and market-interaction layers can be activated or deactivated independently, allowing the incremental contribution of each mechanism to be assessed. This property is central to the simulation methodology adopted in this dissertation, where eight progressive configurations are evaluated to map the operational and economic trade-offs associated with each additional operational or market layer.

The following sections detail the resource-specific formulations that underpin the complete EMS optimisation model presented in Section 3.5.

3.2 Battery Storage Modelling

BESSs are modelled as individual consumer-level resources within the EMS. Each consumer $i \in \mathscr{I}$, where $\mathscr{I}$ denotes the set of consumers, may be equipped with a local storage unit with nominal capacity $\widehat{C}_i^B$, characterised by its charging and discharging efficiencies, SOC limits, initial SOC, and maximum charging and discharging energy per time step. Within the proposed framework, each BESS enables energy to be shifted across time and can support greater use of local PV generation, reduced procurement during more expensive periods, and portfolio-level mFRR reserve provision.

The storage model preserves the physical independence of each consumer. BESS charging, discharging, and SOC variables are therefore defined for each consumer and time period. Let $E_{i,t}^{B,ch}$ and $E_{i,t}^{B,dch}$ denote the charging and discharging energy of consumer i in period t, respectively. The variable $SOC_{i,t}^B$ represents the corresponding SOC, while $\hat{\eta}_i^{B,ch}$ and $\hat{\eta}_i^{B,dch}$ denote the charging and discharging efficiencies. The sets $\mathscr{T}^{opt}$ and $\mathscr{T}^{SOC}$ denote the optimisation horizon and its one-period extension required to represent the SOC reached after the final optimisation period, respectively, as formally defined in Section 3.5.1.

The BESS model follows the formulation of the base tool [133], adapted here to the individual-consumer level. The SOC of each consumer $i \in \mathscr{I}$ evolves over $t \in \mathscr{T}^{opt}$ as:

$$SOC_{i,t+1}^B = SOC_{i,t}^B + \hat{\eta}_i^{B,ch} E_{i,t}^{B,ch} - \frac{1}{\hat{\eta}_i^{B,dch}} E_{i,t}^{B,dch}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.1}$$

Equation 3.1 updates the stored energy by adding the effective charged energy and subtracting the discharged energy corrected by the discharging efficiency. At the beginning of each rolling-horizon cycle, the initial BESS SOC is fixed as:

$$SOC_{i,t_0}^B = \widehat{SOC}_i^{B,init}, \qquad \forall i \in \mathscr{I}. \tag{3.2}$$

In Equation 3.2, t_0 denotes the first period of the current optimisation horizon and $\widehat{SOC}_i^{B,init}$ is the initial SOC of the BESS associated with consumer i. For the first simulated day, this value is obtained from the input data. In subsequent rolling cycles, it is inherited from the end of the previous execution horizon, ensuring intertemporal continuity between consecutive cycles.

The BESS SOC must remain within its admissible technical limits:

$$\widehat{SOC}_i^{B,min} \leq SOC_{i,t}^B \leq \widehat{SOC}_i^{B,max}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{SOC}. \tag{3.3}$$

where $\widehat{SOC}_i^{B,min}$ and $\widehat{SOC}_i^{B,max}$ denote the minimum and maximum admissible SOC values. The upper bound may equal the nominal capacity $\widehat{C}_i^B$, allowing the storage unit to reach its full usable capacity.

Charging and discharging are limited by the maximum energy that can be exchanged in each time step, while simultaneous charging and discharging are prevented through the binary variable $x^B_{i,t}$:

$$0 \leq E^{B,ch}_{i,t} \leq \widehat{E}^{B,ch,max}_{i} \left(1 - x^B_{i,t}\right), \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.4}$$

$$0 \leq E^{B,dch}_{i,t} \leq \widehat{E}^{B,dch,max}_{i} x^B_{i,t}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.5}$$

$$x^B_{i,t} \in \{0,1\}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.6}$$

When $x^B_{i,t} = 0$, the BESS may charge but cannot discharge. When $x^B_{i,t} = 1$, it may discharge but cannot charge. This ensures a physically consistent operating mode in each period.

Although storage operation is modelled individually, the BESS units also support portfolio-level mFRR reserve provision. The available reserve capacity is derived from the charging and discharging headroom and the SOC margins of the individual storage units. The resulting upward and downward reserve bands therefore reflect the combined contribution of all portfolio BESS units rather than the capability of a single storage device. The corresponding reserve-feasibility constraints and aggregation equations are presented in Section 3.5.2.

3.3 Flexible Load Modelling

Flexible loads are represented in the EMS through a DSF modelling layer that derives time-dependent upward and downward flexibility limits from historical consumption data. Unlike approaches that rely on predefined or static flexibility limits, the proposed approach estimates consumer-specific flexibility bounds directly from historical consumption patterns, enabling a more data-driven representation of demand-side flexibility [59], [62]. Instead of assuming fixed or permanently available flexibility margins, the implemented approach estimates the range within which each consumer's demand may be temporarily reshaped according to previously observed consumption behaviour. These estimated margins are then treated as operational bounds in the EMS optimisation model, where the original demand profile can be adjusted through flexibility activation variables.

In the proposed EMS, DSF is interpreted as load shifting rather than permanent demand curtailment. Upward flexibility corresponds to a temporary increase in electricity consumption in a given period, whereas downward flexibility corresponds to a temporary reduction. Reductions must be compensated by equivalent increases within the same execution horizon, preserving the total energy consumption over the executed day while allowing its temporal distribution to be modified.

Historical consumption data are analysed through a six-stage workflow to derive the flexibility bounds used by the optimisation model. Demand data are first preprocessed to obtain a consistent time-indexed structure, and the available history is checked against a minimum length requirement. When sufficient data are available, a configurable historical window is used to estimate the flexibility margins from the selected consumption records. Otherwise, a representative default

profile corresponding to the applicable consumer type and day category is used as a fallback when historical data are insufficient. Fig. 3.2 illustrates the complete workflow.

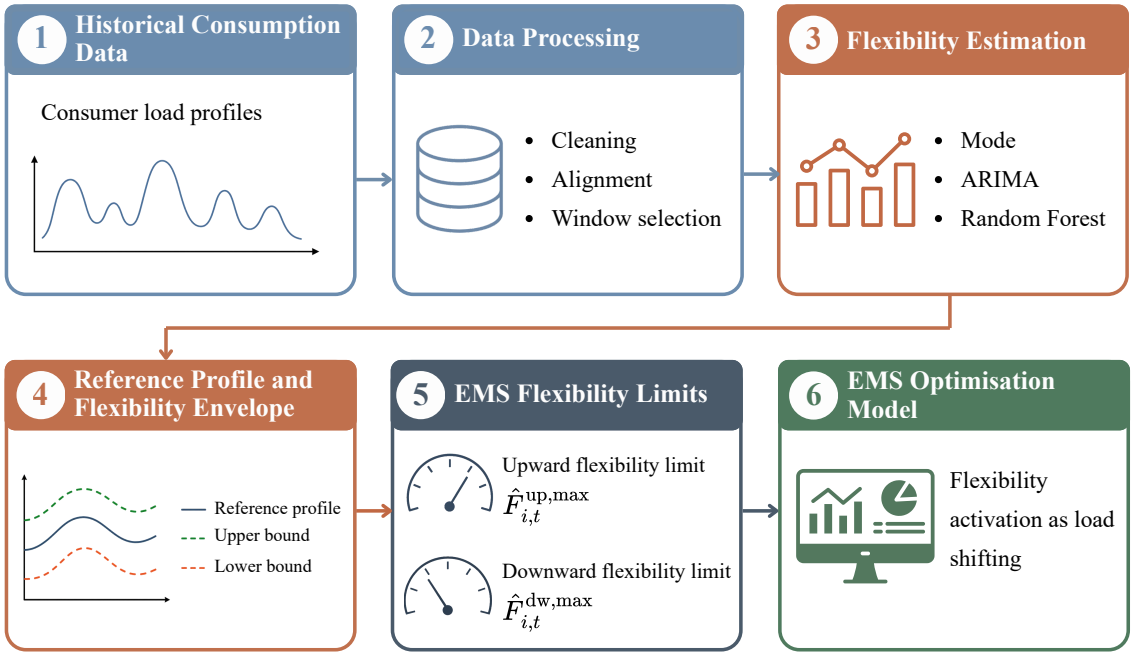


Figure 3.2: Flexible load modelling workflow.

For each consumer $i \in \mathcal{I}$, the module estimates a representative consumption profile together with admissible upper and lower consumption envelopes, consistent with data-driven DSF quantification approaches [59], [134]. The difference between the upper envelope and the reference profile defines the maximum upward flexibility, while the difference between the reference profile and the lower envelope defines the maximum downward flexibility. These margins are converted into the EMS input parameters $\widehat{F}_{i,t}^{up,max}$ and $\widehat{F}_{i,t}^{dw,max}$. The flexibility curves are estimated independently for each weekday and time-of-day slot. When the resulting curves have a 15-minute resolution, the corresponding flexibility margins are aggregated into hourly values and aligned with the simulation calendar before being passed to the optimisation model. The estimated reference profile is used to quantify the flexibility margins and does not replace the consumer load profile used in the EMS energy balance.

The module supports Mode-based, Autoregressive Integrated Moving Average (ARIMA)-based, and Random Forest methods for estimating the reference profile and the flexibility envelope. These alternatives differ in modelling complexity and in the way historical consumption patterns are represented.

The **Mode-based method** follows an empirical approach based on representative consumption profiles, consistent with baseline-oriented DR approaches [134], [135]. For each weekday and time-of-day slot, a representative consumption value is obtained from the historical distribution using a histogram-based mode estimate. The admissible bounds are defined through empirical quantiles at configurable levels and constrained by the historically observed minimum and maxi-

mum values. This approach provides a direct and interpretable representation of typical consumption behaviour without requiring the fitting of a forecasting model.

The **ARIMA-based method** uses a statistical time-series model commonly applied to short-term load forecasting [136]. In the implemented module, the model is fitted to the historical observations associated with each time-of-day slot and used to estimate the corresponding consumption level. The flexibility envelope is derived from the associated forecast interval. This approach can represent temporal dependence in consumption data, although its estimates depend on the characteristics and amount of available historical information.

The **Random Forest method** uses a nonlinear regression model for electricity load forecasting [137]. In the implemented module, temporal features representing the time-of-day, weekday, and seasonal position are used to estimate consumption patterns. The reference profile is obtained from the ensemble prediction, while the flexibility envelope is derived from empirical quantiles of the individual tree predictions, following probabilistic and quantile-based forecasting approaches used in DR applications [138]. This approach can capture nonlinear relationships between temporal features and consumption patterns but requires the fitting of an additional predictive model. Regardless of the selected method, the module produces the same output structure, allowing the resulting margins to be consistently integrated into the optimisation model.

For the simulations conducted in this dissertation, the **Mode-based method** is adopted. The choice reflects the purpose of the flexibility module, which is to derive transparent, consumer-specific operational bounds from observed consumption patterns rather than to optimise short-term load-forecasting accuracy. Its empirical formulation provides a direct link between the historical consumption distribution and the flexibility margins supplied to the EMS, while avoiding additional forecasting-model training and calibration. The ARIMA and Random Forest implementations are retained as alternative estimation approaches. However, an extensive comparison of forecasting accuracy, parameter sensitivity, and out-of-sample robustness among the three methods is outside the scope of this dissertation, whose focus is the operational integration and economic assessment of demand-side flexibility within the aggregated EMS. The adoption of the Mode-based method should therefore be understood as a modelling choice consistent with this objective rather than as evidence of its general superiority over alternative estimators.

In the optimisation model, the upward and downward flexibility of consumer i in period t are represented by $F_{i,t}^{up}$ and $F_{i,t}^{dw}$, respectively. These variables are defined over the optimisation horizon $\mathscr{T}^{opt}$, although flexibility activation is permitted only within the execution horizon $\mathscr{T}^{exec}$, as formally defined in Section 3.5.1. The binary variable $z_{i,t}^{flex}$ enforces mutual exclusivity between directions.

$$0 \leq F_{i,t}^{up} \leq \widehat{F}_{i,t}^{up,max} z_{i,t}^{flex}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{exec}, \tag{3.7}$$

$$0 \leq F_{i,t}^{dw} \leq \widehat{F}_{i,t}^{dw,max} \left(1 - z_{i,t}^{flex}\right), \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{exec}, \tag{3.8}$$

$$z_{i,t}^{flex} \in \{0,1\}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.9}$$

When $z_{i,t}^{flex} = 1$, upward flexibility may be activated and downward flexibility is blocked. When $z_{i,t}^{flex} = 0$, the opposite applies. In either operating mode, the corresponding flexibility variable may remain zero.

Energy neutrality is enforced over the execution horizon for each consumer:

$$\sum_{t \in \mathscr{T}^{exec}} F_{i,t}^{up} = \sum_{t \in \mathscr{T}^{exec}} F_{i,t}^{dw}, \qquad \forall i \in \mathscr{I}. \tag{3.10}$$

Equation 3.10 ensures that the EMS may redistribute consumption within the executed day without changing the consumer's total energy consumption over that horizon. Flexibility is therefore represented as temporary load shifting rather than net demand curtailment.

Flexibility actions are restricted to the execution horizon. No flexibility activation is permitted during the additional look-ahead period, and the corresponding variables are fixed to zero outside $\mathscr{T}^{exec}$:

$$F_{i,t}^{up} = F_{i,t}^{dw} = z_{i,t}^{flex} = 0, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt} \setminus \mathscr{T}^{exec}. \tag{3.11}$$

In the individual energy balance, upward flexibility increases the effective demand of the consumer, whereas downward flexibility reduces it. Load shifting therefore interacts with grid exchanges, BESS operation, EV scheduling, and bilateral procurement, and may indirectly affect the BESS capacity available for reserve provision. DSF does not contribute directly to the portfolio mFRR reserve bands in the proposed formulation. It also has no direct remuneration term in the objective function. Its economic value therefore emerges indirectly through changes in energy procurement and selling, storage operation, reserve availability, and contractual decisions.

3.4 Electric Vehicle Modelling

EVs are modelled in the EMS as flexible consumer-level resources, consistent with aggregator-based scheduling of EV flexibility in prosumer portfolios [114]. Each EV is associated with a specific consumer $i \in \mathscr{I}$. It is characterised by its battery capacity $\widehat{C}_i^{EV}$, charging and discharging efficiencies, SOC limits, arrival SOC, maximum charging and discharging energy per time step, plug-in and plug-out hours, V2G capability, and a required SOC at departure. EV operation is enabled only when a valid plug-in and plug-out window is specified for consumer i. Charging and discharging are explicitly represented through the variables $E_{i,t}^{EV,ch}$ and $E_{i,t}^{EV,dch}$, which enter the individual energy balance and are jointly coordinated with PV generation, BESS operation, DSF, grid exchanges, and bilateral procurement.

A central feature of the EV model is its interaction with the rolling-horizon structure of the EMS, particularly for overnight sessions in which the vehicle connects in the evening and leaves the following morning. Each daily optimisation cycle comprises a 24-hour execution horizon followed by an additional 8-hour look-ahead period. The look-ahead allows the EMS to account for the next-morning plug-out requirement before reaching the daily execution boundary. For the early-morning portion of an overnight session, the relevant EV decisions obtained in the look-ahead are retained and imposed when the same periods reappear in the following optimisation

cycle. This prevents an overnight schedule that has already been used to satisfy the departure requirement from being inconsistently revised. Without this anticipatory treatment, independent day-by-day optimisation could fail to account for a departure requirement occurring after midnight.

Fig. 3.3 illustrates this cross-cycle mechanism, showing the overlap between cycles $D-1$ and D and the transfer of the committed early-morning EV decisions.

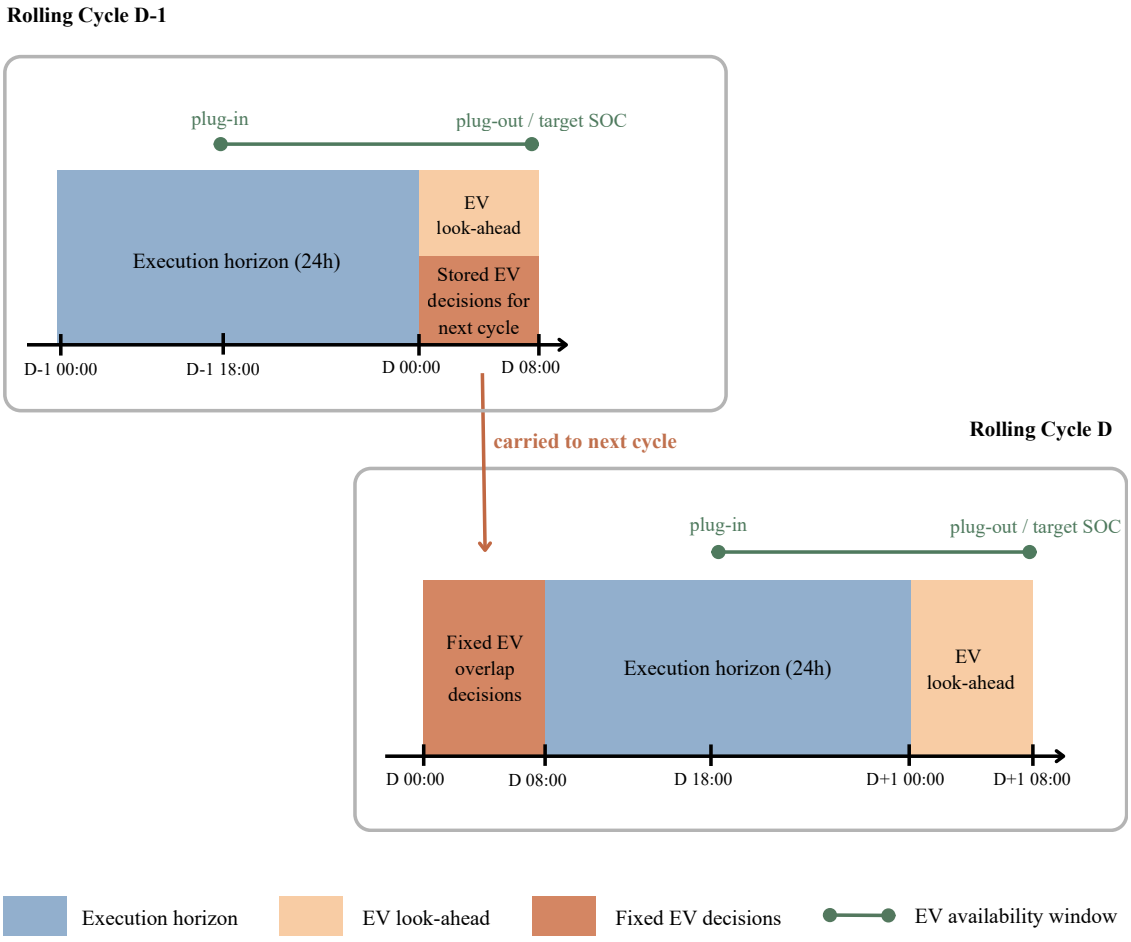


Figure 3.3: Rolling-horizon treatment of an overnight EV session with fixed overlap decisions.

Let $\hat{t}_i^{in}$ and $\hat{t}_i^{out}$ denote the plug-in and plug-out hours of consumer i, and let $h(t)$ denote the hour of the day corresponding to period t. For a same-day availability window ($\hat{t}_i^{in} < \hat{t}_i^{out}$), the EV is available when:

$$\hat{t}_i^{in} \leq h(t) < \hat{t}_i^{out}. \tag{3.12}$$

For an overnight availability window ($\hat{t}_i^{in} > \hat{t}_i^{out}$), the EV is available from the plug-in hour through midnight and from midnight until the plug-out hour of the following day:

$$h(t) \geq \hat{t}_i^{in} \quad \text{or} \quad h(t) < \hat{t}_i^{out}. \tag{3.13}$$

The initial early-morning window is excluded for overnight patterns. Specifically, the periods preceding $\hat{t}_i^{out}$ on the first simulated day are not treated as available, since they would belong to a session that began before the simulation horizon and for which no preceding optimisation cycle exists to provide a carried-over SOC or previously committed decisions.

Let $\mathscr{T}_i^{avail} \subseteq \mathscr{T}^{opt}$ denote the resulting set of available periods for consumer i, after applying the corresponding availability condition and the initial-window exclusion.

Outside the availability window, the EV charging, discharging, and operating-mode variables are forced to zero:

$$E_{i,t}^{EV,ch} = E_{i,t}^{EV,dch} = x_{i,t}^{EV} = 0, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt} \setminus \mathscr{T}_i^{avail}. \tag{3.14}$$

When the EV is unavailable in both the current and following periods, its SOC variable is propagated unchanged, since traction energy while the vehicle is away is not explicitly modelled. This propagation is not imposed when the following period corresponds to a new plug-in event, since the SOC is then reset to the specified arrival value.

The EV SOC is initialised according to the session context. At the first period of a new plug-in event, it is set to the arrival value specified in the input data:

$$SOC_{i,t}^{EV} = \widehat{SOC}_i^{EV,arr}. \tag{3.15}$$

If a rolling cycle begins while an overnight session is already ongoing, the SOC is instead inherited from the end of the previous execution horizon:

$$SOC_{i,t_0}^{EV} = \widehat{SOC}_i^{EV,init}. \tag{3.16}$$

While the EV is connected, its SOC evolves as:

$$SOC_{i,t+1}^{EV} = SOC_{i,t}^{EV} + \hat{\eta}_i^{EV,ch} E_{i,t}^{EV,ch} - \frac{1}{\hat{\eta}_i^{EV,dch}} E_{i,t}^{EV,dch}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}_i^{avail}. \tag{3.17}$$

where $\hat{\eta}_i^{EV,ch}$ and $\hat{\eta}_i^{EV,dch}$ denote the charging and discharging efficiencies of consumer i. The SOC must remain within its admissible technical limits:

$$\widehat{SOC}_i^{EV,min} \leq SOC_{i,t}^{EV} \leq \widehat{SOC}_i^{EV,max}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{SOC}. \tag{3.18}$$

where $\widehat{SOC}_i^{EV,min}$ and $\widehat{SOC}_i^{EV,max}$ denote the minimum and maximum admissible SOC values. The upper bound may equal the nominal capacity $\widehat{C}_i^{EV}$, allowing the vehicle battery to reach its full usable capacity.

At each valid plug-out instant, the EV must satisfy the required departure SOC:

$$SOC_{i,t}^{EV} \geq \widehat{SOC}_i^{EV,target}, \qquad h(t) = \hat{t}_i^{out}. \tag{3.19}$$

For an overnight session, the plug-out instant occurs on the following morning and is included within the look-ahead of the cycle in which the session begins. The departure requirement is therefore enforced when that session is scheduled, before the daily execution boundary. The resulting early-morning decisions are subsequently retained in the following cycle, preserving the charging schedule used to satisfy the required SOC at plug-out.

Charging and discharging are limited by the maximum energy that can be exchanged per time step and cannot occur simultaneously. As with the BESS, these limits are expressed in energy units per time step. Let $x_{i,t}^{EV}$ be the binary operating-mode variable and $\widehat{\delta}_i^{V2G}$ the V2G flag, equal to 1 when V2G is enabled and 0 otherwise:

$$0 \leq E_{i,t}^{EV,ch} \leq \widehat{E}_i^{EV,ch,max}\left(1 - x_{i,t}^{EV}\right), \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{opt}, \tag{3.20}$$

$$0 \leq E_{i,t}^{EV,dch} \leq \widehat{E}_i^{EV,dch,max}\widehat{\delta}_i^{V2G} x_{i,t}^{EV}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{opt}, \tag{3.21}$$

$$x_{i,t}^{EV} \in \{0,1\}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{opt}. \tag{3.22}$$

When $x_{i,t}^{EV} = 0$, the EV may charge but cannot discharge. When $x_{i,t}^{EV} = 1$, the reverse holds. If V2G is disabled ($\widehat{\delta}_i^{V2G} = 0$), discharging is blocked entirely.

The fixed overlap decisions illustrated in Fig. 3.3 are represented through the binary parameter $\widehat{\delta}_{i,t}^{EV,fix}$. A value of 1 indicates that the corresponding EV decision was determined in the look-ahead of the previous cycle and must therefore be retained. The associated charging and discharging energies are denoted by $\widehat{E}_{i,t}^{EV,ch,fix}$ and $\widehat{E}_{i,t}^{EV,dch,fix}$:

$$E_{i,t}^{EV,ch} = \widehat{E}_{i,t}^{EV,ch,fix}, \qquad E_{i,t}^{EV,dch} = \widehat{E}_{i,t}^{EV,dch,fix}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{opt} : \widehat{\delta}_{i,t}^{EV,fix} = 1. \tag{3.23}$$

The fixed values include periods in which zero energy exchange was previously scheduled. Consequently, an interval committed as idle cannot be reconsidered when the same physical period appears in the next optimisation cycle. Together with the carried-over SOC, this mechanism preserves both state and decision continuity across the execution boundary and prevents inconsistent rescheduling of an ongoing overnight session.

3.5 EMS Optimisation Model

3.5.1 Problem Description

The EMS optimisation problem determines the operational schedule that minimises the net variable operating cost of an aggregated consumer portfolio over a rolling horizon. At each optimisation cycle, the model jointly coordinates grid exchanges, local resource operation, demand-side flexibility, reserve provision, and electricity procurement, subject to the technical and operational constraints of each individual consumer.

The formulation follows the aggregation logic introduced in Section 3.1. Each consumer $i \in \mathcal{I}$ is represented through an individual energy balance, grid exchange limits, and the constraints associated with its available BESS, EV, and DSF resources. Decisions requiring portfolio-level coordination, particularly aggregated reserve provision and bilateral procurement limits, are managed collectively by the aggregator. This structure preserves consumer-level feasibility while enabling coordinated operation of the portfolio within a single optimisation problem.

The optimisation is implemented through a rolling-horizon structure. Each daily cycle considers a 32-hour optimisation horizon $\mathcal{T}^{opt}$, comprising a 24-hour execution horizon $\mathcal{T}^{exec}$ and an additional 8-hour look-ahead period. The look-ahead is introduced primarily to anticipate requirements associated with overnight EV sessions, as described in Section 3.4, while extending the optimisation beyond the daily execution boundary. SOC variables for both BESS and EV are defined over the extended set $\mathcal{T}^{SOC}$, which contains one additional period to represent the state resulting from the final optimisation interval. Only decisions within $\mathcal{T}^{exec}$ are treated as executed operational results in each rolling cycle. The states required for intertemporal continuity are transferred to the following cycle, while the relevant early-morning EV overlap decisions determined in the look-ahead are retained and fixed when the corresponding periods are optimised again. Model components that apply specifically to the executed day, including reserve offers and associated revenues, DSF activation, and daily bilateral procurement limits, are therefore restricted to $\mathcal{T}^{exec}$.

The resulting optimisation problem is formulated as a deterministic MILP. Unless otherwise specified, continuous decision variables are non-negative and represent energy exchanges, SOC trajectories, flexibility activation, reserve quantities, and procurement allocations, while binary variables enforce mutually exclusive operating modes and other discrete operational decisions. The complete mathematical formulation is presented in Section 3.5.2, integrating the resource-specific models developed in Sections 3.2, 3.3, and 3.4.

3.5.2 Mathematical Formulation

For the most complete configuration, the objective function is expressed as presented below. Electricity procurement is decomposed into bilateral, market, and tariff channels, while reserve provision and DSF are simultaneously active. The remaining optimised configurations described in Section 4.2 correspond to reduced forms of this formulation, in which specific procurement channels, reserve terms, or DSF variables are omitted or replaced by the corresponding lower-dimensional representation. The objective accounts for electricity procurement costs, revenues from electricity sales, mFRR reserve-band remuneration, and expected activation revenues:

$$\begin{aligned}\min J = &\sum_{t\in\mathcal{T}^{opt}}\sum_{i\in\mathcal{I}}\left(\hat{\pi}^{bil}_{i,t}E^{bil}_{i,t} + \hat{\pi}^{mkt,buy}_{i,t}E^{mkt}_{i,t} + \hat{\pi}^{tar}_{i,t}E^{tar}_{i,t} - \hat{\pi}^{mkt,sell}_{i,t}E^{sell}_{i,t}\right)\\ &+ \sum_{t\in\mathcal{T}^{exec}}\left[-\hat{\pi}^{band}_{t}\left(B^{up}_{t}+B^{dw}_{t}\right) - \hat{\pi}^{r,up}_{t}\rho^{up}_{t}B^{up}_{t} - \hat{\pi}^{r,dw}_{t}\rho^{dw}_{t}B^{dw}_{t}\right].\end{aligned} \tag{3.24}$$

In Equation 3.24, $E^{bil}_{i,t}$, $E^{mkt}_{i,t}$, and $E^{tar}_{i,t}$ represent the electricity purchased by consumer i through the bilateral, market, and tariff channels, respectively, while $E^{sell}_{i,t}$ represents electricity sold to the grid. The parameters $\hat{\pi}^{bil}_{i,t}$, $\hat{\pi}^{mkt,buy}_{i,t}$, $\hat{\pi}^{tar}_{i,t}$, and $\hat{\pi}^{mkt,sell}_{i,t}$ are the corresponding procurement and selling prices. The market channels are referenced to the DAM price series, while the tariff channel reflects the applicable retail tariff structure.

At portfolio level, B^{up}_t and B^{dw}_t represent the aggregated upward and downward mFRR reserve bands. The parameter $\hat{\pi}^{band}_t$ denotes the reserve-band remuneration price, while $\hat{\pi}^{r,up}_t$ and $\hat{\pi}^{r,dw}_t$ denote the upward and downward activation prices. The ratios ρ^{up}_t and ρ^{dw}_t are derived from historical mFRR activation data and are used to estimate the corresponding expected activation revenues, as described in Section 4.1.3.

Procurement costs enter the objective with a positive sign, whereas revenues from electricity sales, reserve-band remuneration, and expected activation enter with a negative sign. Energy-related terms are evaluated over $\mathscr{T}^{opt}$, whereas reserve-related revenues are restricted to $\mathscr{T}^{exec}$ because reserve offers in the look-ahead period are not executed. By integrating procurement costs with revenues from electricity sales, reserve-band remuneration, and expected activation within the same objective function, the formulation captures the economic trade-offs among the available market and procurement strategies.

The objective function is subject to constraints governing individual grid exchanges, local resource operation, DSF, reserve provision, and bilateral procurement, which are presented next by component.

Energy balance and grid exchange constraints: The individual energy balance is imposed at the grid connection point of each consumer. For each consumer $i \in \mathscr{I}$ and period $t \in \mathscr{T}^{opt}$:

$$\begin{aligned} E^{buy}_{i,t} - E^{sell}_{i,t} = &- \widehat{PV}_{i,t} + \widehat{L}_{i,t} + E^{B,ch}_{i,t} - E^{B,dch}_{i,t} \\ &+ E^{EV,ch}_{i,t} - E^{EV,dch}_{i,t} + F^{up}_{i,t} - F^{dw}_{i,t}, \end{aligned} \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{opt}. \tag{3.25}$$

The net grid exchange equals the difference between gross import $E^{buy}_{i,t}$ and gross export $E^{sell}_{i,t}$. In configurations with multiple procurement channels, $E^{buy}_{i,t}$ represents the total purchased energy and is subsequently decomposed into the corresponding bilateral, market, and tariff components. The parameters $\widehat{L}_{i,t}$ and $\widehat{PV}_{i,t}$ denote the deterministic load and PV generation profiles of consumer i, consistent with the perfect-foresight assumption adopted in this work. BESS charging $E^{B,ch}_{i,t}$, EV charging $E^{EV,ch}_{i,t}$, and upward flexibility $F^{up}_{i,t}$ increase the effective demand at the grid connection point. BESS discharging $E^{B,dch}_{i,t}$, EV discharging $E^{EV,dch}_{i,t}$, and downward flexibility $F^{dw}_{i,t}$ reduce it. EV operation is explicitly represented through $E^{EV,ch}_{i,t}$ and $E^{EV,dch}_{i,t}$ and is therefore not included in the baseline load profile $\widehat{L}_{i,t}$. When DSF is not active, $F^{up}_{i,t}$ and $F^{dw}_{i,t}$ are set to zero.

When a contracted power limit is defined for consumer i, gross import and gross export are individually constrained at the grid connection point:

$$E^{buy}_{i,t} \leq \widehat{P}^{contr}_i, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{opt}, \tag{3.26}$$

$$E_{i,t}^{sell} \leq \widehat{P}_i^{contr}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.27}$$

The formulation applies the same contracted-capacity value as a symmetric limit on import and export at the individual grid connection. The input contracted capacity is specified in kVA. Under the unity power-factor assumption adopted in the model, its numerical value is treated as an equivalent active-power limit in kW. Since $\Delta t = 1$ h throughout the formulation, this value also acts directly as the per-period energy limit in kWh for $E_{i,t}^{buy}$ and $E_{i,t}^{sell}$.

Electricity purchase and sale cannot occur simultaneously in any given period. Resource-based upper bounds for the energy that may flow in each direction are used as the Big-M coefficients:

$$\widehat{M}_{i,t}^{buy} = \widehat{L}_{i,t} + \widehat{E}_i^{B,ch,max} + \widehat{E}_i^{EV,ch,max} + \widehat{F}_{i,t}^{up,max}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.28}$$

$$\widehat{M}_{i,t}^{sell} = \widehat{PV}_{i,t} + \widehat{E}_i^{B,dch,max} + \widehat{E}_i^{EV,dch,max}\widehat{\delta}_i^{V2G} + \widehat{F}_{i,t}^{dw,max}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.29}$$

The contracted power constraints in Equations 3.26 and 3.27 may impose tighter limits than these resource-based bounds. When DSF is not active, the flexibility terms in Equations 3.28 and 3.29 are set to zero. Let $y_{i,t}^{grid} \in \{0,1\}$ be the binary variable defining the grid exchange direction. Mutual exclusivity is enforced as:

$$0 \leq E_{i,t}^{buy} \leq \widehat{M}_{i,t}^{buy} y_{i,t}^{grid}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.30}$$

$$0 \leq E_{i,t}^{sell} \leq \widehat{M}_{i,t}^{sell} \left(1 - y_{i,t}^{grid}\right), \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.31}$$

$$y_{i,t}^{grid} \in \{0,1\}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.32}$$

When $y_{i,t}^{grid} = 1$, the consumer may purchase from the grid but cannot sell. When $y_{i,t}^{grid} = 0$, the consumer may sell to the grid but cannot purchase.

Battery storage constraints: The energy-scheduling constraints of each BESS associated with consumer $i \in \mathscr{I}$ comprise the SOC dynamics, initial condition, SOC limits, and charge/discharge exclusivity introduced in Equations 3.1-3.6. These constraints are active over $\mathscr{T}^{opt}$ and, for the SOC limits, over their one-period extension $\mathscr{T}^{SOC}$, ensuring physical feasibility of storage operation for each consumer independently. Beyond energy scheduling, the available charging and discharging headroom and SOC margins provide the physical basis for portfolio-level mFRR reserve provision, as formalised in the reserve constraints below.

Demand-side flexibility constraints: Upward and downward flexibility activation is bounded by the data-driven limits, subject to directional exclusivity, energy neutrality over $\mathscr{T}^{exec}$, and look-ahead blocking, as defined in Equations 3.7-3.11. These constraints ensure that DSF represents within-day load shifting within the estimated flexibility bounds, with no net reduction in total energy demand over the executed day.

Electric vehicle constraints: EV scheduling across consecutive optimisation cycles is governed by the availability conditions, SOC dynamics, technical limits, departure target, and fixed overnight overlap decisions introduced in Equations 3.12-3.23. Charging and discharging are scheduled only while the vehicle is connected. The departure target is enforced when the corresponding session is scheduled, while overnight sessions preserve both state and decision continuity across consecutive rolling cycles.

Reserve provision constraints: The reserve-band formulation is adapted from the base tool and extended here to represent portfolio-level mFRR provision jointly supported by the individual BESS units [133]. Reserve provision is restricted to $\mathscr{T}^{exec}$. Since $\Delta t = 1$ h, the reserve components and aggregated bands are expressed in kWh per period, with numerical values equal to the equivalent average reserve power in kW over the same interval.

For each consumer $i \in \mathscr{I}$ and period $t \in \mathscr{T}^{exec}$, auxiliary variables $S^{ch}_{i,t}$ and $S^{dch}_{i,t}$ represent the unused charging and discharging capability remaining after the scheduled BESS operating point has been determined:

$$E^{B,ch}_{i,t} + S^{ch}_{i,t} = \widehat{E}^{B,ch,max}_{i}, \qquad E^{B,dch}_{i,t} + S^{dch}_{i,t} = \widehat{E}^{B,dch,max}_{i}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{exec}. \tag{3.33}$$

These variables quantify the unused charging and discharging capability available relative to the scheduled BESS operating point. The individual reserve components, $R^{up,ch}_{i,t}$, $R^{up,dch}_{i,t}$, $R^{dw,ch}_{i,t}$, and $R^{dw,dch}_{i,t}$, are bounded by this headroom and by the scheduled operation:

$$0 \le R^{up,dch}_{i,t} \le S^{dch}_{i,t}, \qquad 0 \le R^{up,ch}_{i,t} \le E^{B,ch}_{i,t}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{exec}, \tag{3.34}$$

$$0 \le R^{dw,ch}_{i,t} \le S^{ch}_{i,t}, \qquad 0 \le R^{dw,dch}_{i,t} \le E^{B,dch}_{i,t}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{exec}. \tag{3.35}$$

For upward reserve, $R^{up,dch}_{i,t}$ represents additional discharging capability, while $R^{up,ch}_{i,t}$ represents the ability to reduce scheduled charging. For downward reserve, $R^{dw,ch}_{i,t}$ represents additional charging capability, while $R^{dw,dch}_{i,t}$ represents the ability to reduce scheduled discharging. Together, these four components represent reserve provision as a feasible deviation from the scheduled BESS operating point.

Reserve feasibility must also be guaranteed with respect to the energy stored in each BESS. Upward reserve is limited by the energy available above the minimum SOC, while downward reserve is limited by the remaining capacity below the maximum:

$$\frac{R^{up,dch}_{i,t}}{\hat{\eta}^{B,dch}_{i}} + \hat{\eta}^{B,ch}_{i} R^{up,ch}_{i,t} \le SOC^{B}_{i,t+1} - \widehat{SOC}^{B,min}_{i}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{exec}, \tag{3.36}$$

$$\frac{R^{dw,dch}_{i,t}}{\hat{\eta}^{B,dch}_{i}} + \hat{\eta}^{B,ch}_{i} R^{dw,ch}_{i,t} \le \widehat{SOC}^{B,max}_{i} - SOC^{B}_{i,t+1}, \qquad \forall i \in \mathscr{I},\ \forall t \in \mathscr{T}^{exec}. \tag{3.37}$$

Since reserve is represented as a capacity band rather than as explicitly activated energy, feasibility is evaluated relative to the scheduled end-of-period SOC, $SOC^{B}_{i,t+1}$. The constraints therefore ensure that the offered capacity remains compatible with the available energy margins of each storage unit.

The aggregated portfolio reserve bands are obtained by summing the individual BESS contributions:

$$B^{up}_t = \sum_{i \in \mathscr{I}} \left(R^{up,ch}_{i,t} + R^{up,dch}_{i,t} \right), \qquad B^{dw}_t = \sum_{i \in \mathscr{I}} \left(R^{dw,ch}_{i,t} + R^{dw,dch}_{i,t} \right), \qquad \forall t \in \mathscr{T}^{exec}. \tag{3.38}$$

The formulation therefore represents a single aggregated upward and downward reserve offer at portfolio level, supported by the combined capability of the individual BESS units. B^{up}_t and B^{dw}_t denote reserve-capacity bands rather than physically activated energy. The expected economic effect of reserve activation is accounted for only in the objective function through Equation 3.24.

The reserve-product representation adopted in this work retains the fixed proportionality between the two bands used in the source formulation, under which the upward band is twice the downward band [114]:

$$B^{up}_t = 2\, B^{dw}_t, \qquad \forall t \in \mathscr{T}^{exec}. \tag{3.39}$$

This 2:1 relation is inherited from the historical MIBEL reserve-band representation used in that formulation, in which the upward and downward components corresponded to two-thirds and one-third of the total band, respectively. It is retained as a modelling assumption for consistency with the inherited formulation and the case-study parameterisation, and should not be interpreted as a general or current mFRR market requirement. The proportionality factor can be modified without requiring structural changes to the formulation.

Reserve participation is represented by the binary variable u^{res}_t. To link the offered bands to this participation decision, the aggregate Big-M coefficient is defined as:

$$\widehat{M}^{res} = \sum_{i \in \mathscr{I}} \left(\widehat{E}^{B,ch,max}_i + \widehat{E}^{B,dch,max}_i \right). \tag{3.40}$$

The parameter $\widehat{M}^{res}$ provides a conservative upper bound on the aggregated upward reserve band based on the maximum charging and discharging capability of all portfolio BESS units per time step, rather than on their nominal energy capacity. Given the fixed relation in Equation 3.39, $\widehat{M}^{res}/2$ provides a tighter upper bound for the downward band. The reserve bands are linked to the participation decision as:

$$B^{up}_t \leq \widehat{M}^{res}\, u^{res}_t, \qquad B^{dw}_t \leq \frac{\widehat{M}^{res}}{2}\, u^{res}_t, \qquad \forall t \in \mathscr{T}^{exec}. \tag{3.41}$$

$$u^{res}_t \in \{0,1\}, \qquad \forall t \in \mathscr{T}^{opt}. \tag{3.42}$$

A minimum participation threshold $\widehat{P}^{res,min}$ is imposed on the downward band. Equation 3.39 sets the corresponding upward minimum, so no separate constraint is needed:

$$B_t^{dw} \geq \widehat{P}^{res,min} u_t^{res}, \qquad \forall t \in \mathscr{T}^{exec}. \tag{3.43}$$

The parameter $\widehat{P}^{res,min}$ is a configurable modelling threshold whose numerical value is specified at case-study level, as discussed in Section 4.1.4. It does not represent the actual minimum mFRR market participation requirement. When $u_t^{res} = 0$, no reserve band is offered. When $u_t^{res} = 1$, the portfolio must offer at least $\widehat{P}^{res,min}$ in the downward band and, through the fixed proportionality constraint, at least $2\widehat{P}^{res,min}$ in the upward band.

All reserve variables are fixed to zero outside the execution horizon:

$$B_t^{up} = B_t^{dw} = u_t^{res} = 0, \qquad \forall t \in \mathscr{T}^{opt} \setminus \mathscr{T}^{exec}. \tag{3.44}$$

$$R_{i,t}^{up,ch} = R_{i,t}^{up,dch} = R_{i,t}^{dw,ch} = R_{i,t}^{dw,dch} = S_{i,t}^{ch} = S_{i,t}^{dch} = 0, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt} \setminus \mathscr{T}^{exec}. \tag{3.45}$$

Together, these constraints confine reserve provision to the current execution horizon, while the additional look-ahead period remains available for the underlying energy scheduling problem.

Bilateral procurement constraints: When bilateral procurement is active, the total electricity purchased by each consumer is decomposed into three channels:

$$E_{i,t}^{buy} = E_{i,t}^{bil} + E_{i,t}^{mkt} + E_{i,t}^{tar}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.46}$$

The variables $E_{i,t}^{bil}$, $E_{i,t}^{mkt}$, and $E_{i,t}^{tar}$ denote the energy procured through the bilateral contract, the DAM, and the retail tariff, respectively. This decomposition preserves the individual energy balance in Equation 3.25 while allowing bilateral procurement to be coordinated at portfolio level.

The aggregated bilateral procurement is defined as the sum of the individual allocations across all consumers:

$$E_t^{bil,agg} = \sum_{i \in \mathscr{I}} E_{i,t}^{bil}, \qquad \forall t \in \mathscr{T}^{opt}. \tag{3.47}$$

The aggregate variable $E_t^{bil,agg}$ is used to enforce the contractual limits of the bilateral procurement channel. These comprise an hourly volume cap at portfolio level and minimum and maximum daily procurement requirements:

$$0 \leq E_t^{bil,agg} \leq \widehat{E}^{bil,h}, \qquad \forall t \in \mathscr{T}^{opt}, \tag{3.48}$$

$$\sum_{t \in \mathscr{T}^{exec}} E_t^{bil,agg} \geq \widehat{E}^{bil,min,day}, \tag{3.49}$$

$$\sum_{t \in \mathscr{T}^{exec}} E_t^{bil,agg} \leq \widehat{E}^{bil,max,day}. \tag{3.50}$$

The parameters $\widehat{E}^{bil,h}$, $\widehat{E}^{bil,min,day}$, and $\widehat{E}^{bil,max,day}$ denote the hourly, minimum daily, and maximum daily bilateral procurement limits, respectively. The daily minimum represents the minimum purchase obligation of the modelled take-or-pay structure, while the hourly and daily upper bounds limit the volume procured through the contract. The hourly cap applies throughout $\mathscr{T}^{opt}$ whenever bilateral energy is scheduled, whereas the daily minimum and maximum are evaluated only over $\mathscr{T}^{exec}$. The look-ahead period therefore cannot contribute to satisfying the contractual obligations of the current executed day.

For the residual non-bilateral purchase, the EMS may use either the DAM or the retail tariff channel in each period. These two channels are mutually exclusive, while the objective function determines the economically preferred procurement decision. The binary variable $z^{buy}_{i,t}$ governs this selection:

$$0 \leq E^{mkt}_{i,t} \leq \widehat{M}^{buy}_{i,t} z^{buy}_{i,t}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.51}$$

$$0 \leq E^{tar}_{i,t} \leq \widehat{M}^{buy}_{i,t} \left(1 - z^{buy}_{i,t}\right), \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}, \tag{3.52}$$

$$z^{buy}_{i,t} \in \{0,1\}, \qquad \forall i \in \mathscr{I}, \forall t \in \mathscr{T}^{opt}. \tag{3.53}$$

When $z^{buy}_{i,t} = 1$, procurement through the DAM is permitted and the retail tariff channel is blocked. When $z^{buy}_{i,t} = 0$, the retail tariff channel is permitted and the DAM is blocked. In either case, the permitted variable may remain zero. Bilateral procurement is modelled independently through $E^{bil}_{i,t}$ and may therefore coexist in the same period with either permitted residual channel.

3.6 Computational Implementation and Result Processing

The EMS framework described in this chapter was implemented in Python using the Pyomo optimisation modelling environment and solved with the COIN-OR Branch and Cut (CBC) open-source MILP solver. The computational workflow comprises input and configuration processing, model construction, rolling-horizon execution, result storage, and post-processing. This workflow connects the mathematical formulation to the operational and economic outputs used in the subsequent analysis.

The simulations are coordinated by a main execution routine that defines the portfolio composition, simulation period, active modelling layers, and solver configuration. For each daily cycle of an optimised configuration, the Pyomo model is rebuilt using the corresponding input data and the state information inherited from the preceding cycle. The complete optimisation horizon is solved, after which only the results corresponding to the execution period are retained as executed operational outputs. The state values required for intertemporal continuity are then transferred to the next cycle. For overnight EV sessions, the relevant overlap decisions determined in the look-ahead are additionally preserved when the corresponding early-morning periods reappear in the subsequent optimisation cycle.

The solver is operated through a configurable profile specifying the maximum solution time and the relative and absolute Mixed-Integer Programming (MIP) gap tolerances. When a solution

from the preceding cycle is available, a time-shifted version is used to initialise the following optimisation, providing an informed starting point without altering the underlying mathematical formulation. Solver status and termination information are retained throughout the simulation workflow to support the assessment of computational performance and solution quality.

The baseline configurations defined in Section 4.2 are evaluated separately as reference operating configurations, while each optimised configuration is executed as an independent sequence of rolling-horizon cycles. SOC values and committed overnight EV decisions are carried only between consecutive cycles belonging to the same configuration. Before a new configuration is started, these states and stored decisions are reset to their original input conditions, ensuring that the outcome of one simulated configuration does not influence the initial conditions of another.

The executed results are stored in structured intermediate data and subsequently consolidated into portfolio-level and individual-consumer Excel reports. These outputs contain the operational schedules and economic quantities required for the comparative analysis, including grid exchanges, resource operation, procurement by channel, reserve provision, flexibility activation, and the corresponding cost and revenue components. At individual level, the portfolio-level reserve-band remuneration is allocated during post-processing proportionally to the contribution of each consumer's BESS to the aggregated reserve bands. This allocation is applied only to individual economic reporting and does not alter the optimised dispatch, the aggregated reserve provision, or the portfolio-level economic results.

The processed outputs also support consistency and solution-quality checks across the simulated configurations, including energy balances, demand-flexibility neutrality, reserve-band relationships, bilateral contractual limits, EV departure requirements, and solver termination status. These outputs form the numerical basis for the comparative assessment in Chapter 4, where the progressive operational and market configurations are evaluated in terms of their technical and economic performance.

Chapter 4

Case Study and Trade-Off Analysis

This chapter applies the aggregated EMS developed in Chapter 3 to a B2C consumer portfolio operating within the Portuguese electricity market context. The case study is structured as a progressive evaluation of eight EMS configurations, from reference baselines to a fully active multi-market configuration combining storage scheduling, EV management, reserve provision, DSF, and bilateral procurement, allowing the incremental impact of each mechanism and the resulting trade-offs to be systematically assessed.

The chapter begins with a description of the consumer portfolio, asset parameters, market price inputs, and simulation setup, followed by the definition of the eight evaluated configurations. Results are then presented and discussed at portfolio and individual-consumer levels, including a seasonal analysis, before concluding with a synthesis of the main trade-offs identified across configurations.

4.1 Case Study Description and Simulation Setup

The proposed aggregated EMS is evaluated over the full calendar year 2022 using a B2C consumer portfolio operating within the Portuguese electricity market framework. The case study was developed within the scope of the ATE research project and draws on real consumption and PV generation data collected from residential consumers monitored at SEL. All input data used in the case study, together with the resulting portfolio-level and individual-consumer-level result reports, are publicly available in the reference dataset [139].

Raw consumption and generation data are provided at 15-minute resolution and aggregated to the hourly time step ($\Delta t = 1$ h) used throughout the EMS formulation (Section 3.5.2). This temporal resolution is also consistent with the hourly market price data described in Section 4.1.3. All input profiles, including consumption, PV generation, and market price series, correspond to historical 2022 realisations and are used directly as deterministic inputs, so that the EMS operates under a perfect-foresight assumption.

4.1.1 Consumer Portfolio and Tariff Scenarios

The portfolio comprises five residential B2C consumers, labelled User 1 to User 5, all connected to the low-voltage network under the Normal Low Voltage (BTN) tariff structure. Each consumer is assigned an ERSE-regulated retail tariff structure [140], which defines the retail tariff price $\hat{\pi}_{i,t}^{tar}$ according to the applicable tariff option, period, and day type. The assignment of tariff cycle and bilateral procurement price is configurable within the simulation framework, allowing different scenarios to be evaluated without changes to the optimisation model. Table 4.1 presents the tariff levels applicable to each consumer.

Table 4.1: BTN retail tariff levels ($\hat{\pi}_{i,t}^{tar}$) by time-of-use cycle [140].

Consumers	**Cycle**	**Off-peak [€/kWh]**	**Full-hour [€/kWh]**	**Peak [€/kWh]**
Users 2, 4	Daily simple (flat)	0.1220	-	-
User 5	Weekly bi-hourly	0.1369	-	0.2095
Users 1, 3	Weekly tri-hourly	0.1072	0.1721	0.3170

On workdays, the daily simple cycle applies a single flat rate across all hours, the weekly bi-hourly cycle distinguishes a daytime peak period (07:00-24:00) from an overnight off-peak, and the weekly tri-hourly cycle adds an intermediate full-hour level between the peak windows (09:00-12:00 and 18:00-21:00) and the off-peak. Both time-differentiated cycles follow separate, ERSE-regulated schedules at weekends, with Sundays priced entirely at the off-peak level.

All five consumers are assigned a bilateral procurement price $\hat{\pi}_{i,t}^{bil} = 0.1369$ €/kWh, which coincides with the off-peak level of the weekly bi-hourly cycle and lies below the 2022 annual average DAM price (0.1679 €/kWh) and below the peak levels of both time-differentiated cycles. This positioning is central to the procurement trade-offs evaluated across the EMS configurations.

Each consumer additionally incurs a daily fixed contracted power charge $\hat{\pi}_{i}^{cp}$, determined by the contracted power $\widehat{P}_{i}^{contr}$ according to the ERSE-regulated standing charge schedule [140]. These costs are invariant across EMS configurations and excluded from the optimisation objective, but are still reported. Table 4.2 summarises $\widehat{P}_{i}^{contr}$ and $\hat{\pi}_{i}^{cp}$ for each consumer.

Table 4.2: Contracted power ($\widehat{P}_{i}^{contr}$) and daily standing charge ($\hat{\pi}_{i}^{cp}$) per consumer [140].

Consumer	$\widehat{P}_{i}^{contr}$ **[kVA]**	$\hat{\pi}_{i}^{cp}$ **[€/day]**
User 1	6.90	0.2567
User 2	3.45	0.1283
User 3	13.80	0.5134
User 4	6.90	0.2567
User 5	4.60	0.1711
Portfolio	**-**	**1.3262**

4.1.2 Distributed Energy Resources

Battery Energy Storage Systems: Each of the five consumers is equipped with a behind-the-meter BESS with a nominal capacity $\widehat{C}_i^B = 5$ kWh. The charge and discharge efficiencies $\hat{\eta}_i^{B,ch} = \hat{\eta}_i^{B,dch} = 0.90$ are uniform across all units, yielding a round-trip efficiency of 81%. The maximum charge and discharge limits $\widehat{E}_i^{B,ch,max} = \widehat{E}_i^{B,dch,max} = 1$ kWh per time step are applied consistently for all consumers. Since $\Delta t = 1$ h, these limits are numerically equivalent to a power rating of 1 kW. The SOC is bounded between $\widehat{SOC}_i^{B,min} = 10\%$ and $\widehat{SOC}_i^{B,max} = 100\%$ of $\widehat{C}_i^B$. Since the upper bound equals the nominal capacity, $\widehat{SOC}_i^{B,max} = \widehat{C}_i^B$ in this case study. At the start of each configuration, the initial SOC is set to $\widehat{SOC}_i^{B,init} = 50\%$ of $\widehat{C}_i^B$. Table 4.3 summarises these parameters, which are uniform across all portfolio consumers.

Table 4.3: BESS parameters applied to all portfolio consumers.

Parameter	Value
Nominal capacity ($\widehat{C}_i^B$)	5.0 kWh
Maximum charge limit ($\widehat{E}_i^{B,ch,max}$)	1.0 kWh
Maximum discharge limit ($\widehat{E}_i^{B,dch,max}$)	1.0 kWh
Charge efficiency ($\hat{\eta}_i^{B,ch}$)	0.90
Discharge efficiency ($\hat{\eta}_i^{B,dch}$)	0.90
Minimum SOC ($\widehat{SOC}_i^{B,min}$)	10%
Maximum SOC ($\widehat{SOC}_i^{B,max}$)	100%
Initial SOC ($\widehat{SOC}_i^{B,init}$)	50%

Photovoltaic Generation: All consumers have local PV generation, with annual outputs ranging from 682 kWh to 4,201 kWh across the portfolio. The PV profiles enter the EMS as the parameter $\widehat{PV}_{i,t}$ and are obtained by aggregating the measured 15-minute generation data to hourly resolution. No forecasting error is introduced, consistent with the perfect-foresight assumption adopted in this case study.

Electric Vehicles: Four of the five consumers own an EV (Users 1, 3, 4, and 5). User 2 has no EV. All EV-equipped consumers follow an overnight availability pattern, with plug-in hours $\hat{t}_i^{in}$ between 16:00 and 18:00 and a common plug-out hour $\hat{t}_i^{out} = 08{:}00$. The charge and discharge efficiencies are uniform across all vehicles, with $\hat{\eta}_i^{EV,ch} = \hat{\eta}_i^{EV,dch} = 0.95$, and the maximum discharge limit equals the maximum charge limit, $\widehat{E}_i^{EV,dch,max} = \widehat{E}_i^{EV,ch,max}$, for all EV-equipped consumers. As with the BESS, these limits are expressed in kWh per time step and are numerically equivalent to a power rating in kW given $\Delta t = 1$ h. The EMS is required to reach a target SOC $\widehat{SOC}_i^{EV,target} = 80\%$ of battery capacity at each plug-out instant. V2G capability is controlled by the flag $\widehat{\delta}_i^{V2G}$, which is set to zero in this case study so that all EVs operate in unidirectional smart-charging mode. Table 4.4 summarises the individual EV parameters.

Table 4.4: EV parameters per consumer.

Parameter	**User 1**	**User 2**	**User 3**	**User 4**	**User 5**
Battery capacity ($\widehat{C}_i^{EV}$) [kWh]	40	-	50	40	30
Maximum charge limit ($\widehat{E}_i^{EV,ch,max}$) [kWh]	3.7	-	7.4	3.7	2.3
Plug-in hour ($\hat{t}_i^{in}$)	17:00	-	18:00	17:00	16:00
Arrival SOC ($\widehat{SOC}_i^{EV,arr}$) [%]	40	-	35	50	30
Target SOC ($\widehat{SOC}_i^{EV,target}$) [%]	80	-	80	80	80

The overnight nature of the EV sessions introduces an intertemporal dependency across consecutive rolling-horizon cycles. As described in Section 3.4, the EMS handles this through an 8-hour look-ahead period that anticipates the plug-out requirement at $\hat{t}_i^{out}$, and through a fixed-decision carry-over mechanism that commits the early-morning charging schedule of ongoing overnight sessions.

Demand-Side Flexibility: Consumer-level DSF limits are estimated from the 2022 historical consumption data using the Mode-based method described in Section 3.3, with a 90-day historical data window and empirical quantiles at the 10th and 90th percentile levels. The resulting upward and downward flexibility bounds, $\widehat{F}_{i,t}^{up,max}$ and $\widehat{F}_{i,t}^{dw,max}$, are computed independently for each weekday and hour of the day and aligned with the simulation calendar before being passed to the optimisation model. The flexibility limits are computed once, ahead of the EMS configurations that use them, and are included only in configurations with explicit DSF enabled, as presented in Section 4.2.

4.1.3 Energy and Reserve Market Price Inputs

Day-Ahead Energy Market: Hourly DAM prices for the Portuguese bidding zone are sourced from OMIE and retrieved from the REN Data Hub [141], covering the full calendar year 2022. The annual average price is 167.9 €/MWh, reflecting the elevated and volatile price environment associated with the 2022 European energy crisis. In the EMS formulation, the effective market buy price $\hat{\pi}_{i,t}^{mkt,buy}$ and sell price $\hat{\pi}_{i,t}^{mkt,sell}$ are both referenced to the same hourly DAM price series, as defined in Section 3.5.2. This price series is available in the reference dataset [139].

Reserve Market: mFRR reserve inputs are derived from historical Portuguese balancing market data for 2022, retrieved from the REN Data Hub [141] and available in the reference dataset [139]. Three price signals enter the EMS objective function directly: the mFRR capacity band price $\hat{\pi}_t^{band}$, which remunerates contracted reserve capacity, and the upward and downward activation prices $\hat{\pi}_t^{r,up}$ and $\hat{\pi}_t^{r,dw}$, which reflect the marginal price of activated mFRR energy in each direction, as defined in Section 3.5.2.

The expected activation ratios ρ_t^{up} and ρ_t^{dw}, also used in the objective function to estimate activation revenues, are not retrieved directly but computed from the same source. For each hour t,

they are defined as the ratio of historically activated mFRR energy to the contracted system band:

$$\rho_t^{up} = \frac{E_t^{act,up}}{B_t^{up,sys}\Delta t}, \qquad \rho_t^{dw} = \frac{E_t^{act,dw}}{B_t^{dw,sys}\Delta t}. \tag{4.1}$$

where $E_t^{act,up}$ and $E_t^{act,dw}$ are the upward and downward mFRR energy activated in the Portuguese system in hour t, and $B_t^{up,sys}$ and $B_t^{dw,sys}$ are the corresponding contracted mFRR capacity bands.

Prior to the simulation, an optimistic pre-processing step is applied to the raw ratios: in each hour t, only the direction with the higher ratio is retained and the other is set to zero, ensuring that expected activation revenues are credited in at most one direction per hour. These ratios are used as deterministic scaling factors for the expected activation component of the objective function and should not be interpreted as activation probabilities or as physically simulated activation trajectories.

4.1.4 Computational Setup

Each optimised configuration is executed over the full calendar year 2022, comprising 365 daily rolling-horizon cycles. Each cycle optimises a 32-hour window composed of a 24-hour execution horizon ($\mathscr{T}^{exec}$) and an 8-hour EV look-ahead period, as defined in Section 3.5.2. Only the execution horizon is stored and exported as final output, while the relevant look-ahead decisions are retained to preserve overnight EV scheduling consistency across consecutive cycles. The optimisation model is solved using the CBC open-source MILP solver. Each daily solve is subject to a maximum solution time of 6000 seconds, a relative MIP gap tolerance of 20%, and an absolute MIP gap tolerance of 0.05 €. These tolerances define alternative stopping criteria and do not represent the optimality gaps necessarily reached by the final solutions. In practice, all annual solves terminated successfully under the configured solver criteria for all optimised configurations.

The minimum reserve participation threshold $\widehat{P}^{res,min}$, introduced in Equation 3.43, is set to 0.50 kWh per hourly time step to avoid residual or negligible reserve offers. As discussed in Section 2.3.2, current Portuguese mFRR market rules require a minimum offer of 1 MW, which exceeds the physical capability of the B2C portfolio considered here. The reserve analysis should therefore be interpreted as an assessment of the economic potential of BESS-based reserve provision under a relaxed participation threshold rather than as direct market eligibility of the five-consumer portfolio. Aggregation of sub-1 MW units remains subject to the applicable qualification and technical requirements [92].

The bilateral procurement limits are set to $\widehat{E}^{bil,h} = 2.25$ kWh/h, $\widehat{E}^{bil,min,day} = 3.5$ kWh/day, and $\widehat{E}^{bil,max,day} = 20.0$ kWh/day, as defined in Equations 3.48-3.50. These limits were calibrated from the aggregated PV availability of the portfolio using a conservative scaling factor $\gamma = 0.35$ applied to the 95th percentile of the hourly and daily aggregated PV generation. This calibration defines representative contractual volume limits without imposing a direct physical coupling between bilateral procurement and contemporaneous PV generation.

4.2 Progressive EMS Configurations

The case study evaluates eight EMS configurations: two reference baselines and six progressively enriched optimisation configurations. The configurations are structured to introduce successive operational and market layers, allowing the incremental contribution of each mechanism and the resulting trade-offs to be assessed systematically. Table 4.5 provides an overview of the active layers in each configuration.

Table 4.5: Overview of EMS configurations and active operational layers.

Configuration	Buy channel	Sell	EV	BESS	Reserve	DSF	Bilateral
Baseline Simple	Tariff	-	✓	-	-	-	-
Baseline BESS	Tariff	-	✓	✓	-	-	-
Case A	Tariff	-	✓	✓	-	-	-
Case B	Tariff	Mkt.	✓	✓	-	-	-
Case C	Tar./Mkt.	Mkt.	✓	✓	-	-	-
Case D	Tar./Mkt.	Mkt.	✓	✓	✓	-	-
Case E	Tar./Mkt.	Mkt.	✓	✓	✓	✓	-
Case F	Bil.+Tar./Mkt.	Mkt.	✓	✓	✓	✓	✓

Notes: ✓ = layer active; - = layer inactive.

The two baselines serve as cost references. **Baseline Simple** represents a fully passive portfolio in which all grid purchases are settled at the retail tariff, EVs charge immediately after plug-in at maximum power without scheduling, and no BESS is operated. **Baseline BESS** extends this reference with a rule-based storage heuristic that charges during PV surplus or low-tariff periods and discharges at peak tariff, while EV charging remains uncontrolled. Neither baseline has a remunerated selling channel, so any residual PV surplus is exported without remuneration.

The six optimisation configurations build progressively on each other. **Case A** jointly optimises BESS scheduling and EV smart charging under the retail tariff. **Case B** introduces remunerated grid exports through the aggregator at the DAM price, creating the first trade-off between local energy use and market-based selling. **Case C** enables DAM purchases, allowing the optimiser to choose between the retail tariff and the DAM in each period.

From Case D onwards, additional portfolio-level mechanisms are introduced. **Case D** adds mFRR reserve provision, aggregating the individual BESS capabilities into upward and downward reserve bands remunerated through band and expected activation revenues, at the cost of reduced scheduling freedom. **Case E** introduces DSF, enabling temporary load shifts to be coordinated with BESS scheduling and reserve provision subject to daily energy neutrality. Finally, **Case F** activates bilateral procurement, decomposing grid purchases among the bilateral contract, the DAM, and the retail tariff subject to hourly and daily contractual limits, completing the multi-market configuration evaluated in this dissertation.

4.3 Portfolio-Level Results

4.3.1 Operating Cost Across Cases

The annual variable operating cost of the portfolio provides the first and most direct measure of the value added by each successive modelling layer. Table 4.6 reports the total variable cost of the eight configurations over the full 2022 simulation, together with the savings relative to the fully passive Baseline Simple. This variable cost is the quantity minimised by the EMS in Cases A to F, whereas the two baselines follow the passive and rule-based procedures introduced in Section 4.2. A fixed contracted-power charge of 484.06 €/year applies to the portfolio in all configurations and is excluded from the variable-cost values reported here. The corresponding daily charges are presented in Section 4.1.1.

Table 4.6: Annual variable operating cost and savings relative to Baseline Simple across cases.

Configuration	Variable cost [€/year]	Savings vs. Simple [%]
Baseline Simple	9584.4	-
Baseline BESS	8913.6	7.0
Case A	6165.2	35.7
Case B	5646.9	41.1
Case C	4988.4	47.9
Case D	622.2	93.5
Case E	450.3	95.3
Case F	284.5	97.0

Fig. 4.1 illustrates the main incremental contributions underlying this cost progression.

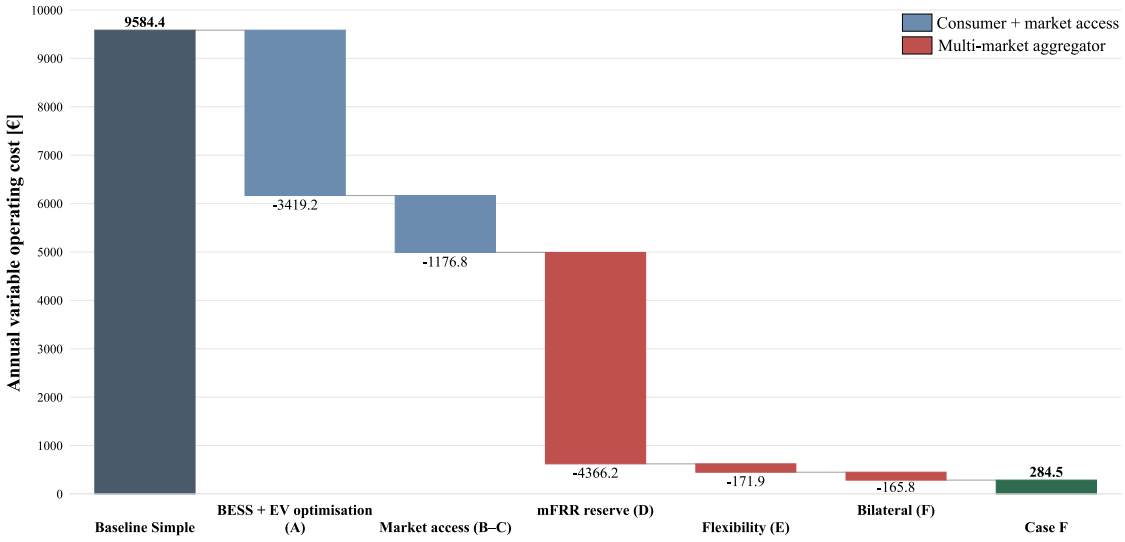


Figure 4.1: Waterfall decomposition of the annual variable operating cost across cases.

The results in Table 4.6 and Fig. 4.1 are organised around two questions that guide the analysis. The **first question** concerns how much economic value can be unlocked by coordinating existing behind-the-meter resources and enabling access to energy-market opportunities through an aggregator, addressed by the progression from the baselines to Case C. The **second question** concerns how much further value can be created by extending this framework to reserve participation, DSF, and bilateral procurement, addressed by Cases D to F.

For the first question, the progression is unambiguous. Operating the available storage under the simple rule-based heuristic of Baseline BESS reduces the cost by only 7.0% relative to the passive reference. The decisive step is instead the coordinated optimisation of storage and EV charging in Case A, which lowers the cost by 35.7% relative to Baseline Simple and by 30.8% relative to Baseline BESS. This gap shows that coordinated scheduling extracts substantially more value from the available assets than rule-based storage operation combined with uncontrolled EV charging. Adding market access through remunerated exports in Case B and market-based purchasing in Case C raises the cumulative saving to 47.9%, nearly halving the portfolio's variable operating cost before reserve provision, DSF, or bilateral procurement are introduced.

For the second question, the additional benefit is considerably larger than the market-access step that precedes it. Introducing mFRR reserve provision in Case D lowers the cost from 4988.4 € to 622.2 €, a single step worth 4366.2 € that raises the cumulative saving to 93.5%. This is visible in Fig. 4.1 as the largest incremental reduction across all cases, and under the assumptions adopted in this case study, reserve provision emerges as the main economic driver of the multi-market configuration. The two remaining layers operate in a regime of diminishing returns: DSF in Case E reduces the cost by a further 171.9 €, while bilateral procurement in Case F contributes an additional 165.8 €, bringing the fully active configuration to 284.5 € and a cumulative saving of 97.0%. From Case E onwards, the remaining variable cost falls below the fixed contracted-power charge of 484.06 € per year, showing how small the remaining net variable operating cost becomes after the successive optimisation and market layers have been introduced.

The following subsections examine the mechanisms behind this cost progression, beginning with the energy scheduling, procurement, and selling decisions that distinguish Cases A to C, before examining the storage and reserve operation behind the dominant step in Case D, EV scheduling, and the subsequent contributions of DSF and bilateral procurement.

4.3.2 Energy Procurement and Selling

The savings observed across Cases A to C are driven primarily by how the portfolio operates its available resources and exchanges energy with the grid, rather than by any material change in its underlying energy needs. With a fixed annual demand of 33,698 kWh, an EV charging requirement of around 25,100 kWh that varies only marginally across configurations, and 12,551 kWh of local PV generation, the portfolio remains a strong net importer throughout the case study. The two baselines provide passive and rule-based references under retail-tariff procurement, while Cases A to C progressively introduce coordinated BESS and EV scheduling, remunerated grid exports, and DAM procurement through the aggregated EMS. In these optimised cases, the five consumers are

coordinated jointly while each retains its individual energy balance and grid connection. Table 4.7 reports the resulting annual grid exchanges from Baseline Simple to Case C. Since gross import and export are summed across the five individual connection points rather than measured at a single aggregated meter, simultaneous imports and exports by different consumers contribute to both gross totals. Their difference gives the portfolio net import and is unaffected by this reporting distinction.

Table 4.7: Annual portfolio grid exchanges from Baseline Simple to Case C.

Configuration	**Gross import [kWh]**	**Gross export [kWh]**	**Net import [kWh]**
Baseline Simple	51,354	5,041	46,313
Baseline BESS	50,303	3,036	47,267
Case A	49,116	1,672	47,444
Case B	55,210	7,345	47,866
Case C	57,444	9,314	48,130

Case A optimises storage and EV charging under the retail tariff but has no remunerated selling channel, so the BESS is scheduled purely to reduce operating cost. Both gross flows fall relative to Baseline Simple. Battery discharge displaces grid imports that would otherwise be required, while charging absorbs PV surplus that would otherwise be exported, reducing annual import from 51,354 to 49,116 kWh and export from 5,041 to 1,672 kWh, the lowest export level across all configurations. This residual export represents instantaneous surplus that cannot be absorbed within the same hour rather than curtailed generation. Much of this residual export occurs when the 1 kWh/h charging limit prevents further absorption despite remaining storage capacity. Net import nevertheless rises by about 2.4% relative to Baseline Simple because energy cycled through the BESS is subject to its 81% round-trip efficiency, and the associated losses must ultimately be supplied through additional grid imports.

Opening the remunerated selling channel in Case B and market-based purchasing in Case C changes this operating pattern substantially. By Case C, gross import reaches 57,444 kWh and gross export 9,314 kWh, while net import rises by only about 1.5% relative to Case A. Two effects explain this divergence. First, greater use of storage for price-driven arbitrage increases the energy cycled through the BESS and therefore the associated round-trip losses. More importantly, in Case C, some consumers export energy while others simultaneously import in more than half of the hours. Because exchanges are recorded at five separate connection points, these opposing transactions increase both gross totals even though they partially offset when viewed at portfolio level. Market access therefore does not materially reduce the portfolio's underlying energy requirement. Instead, it increases the volume of economically driven exchanges around an essentially fixed physical need.

Even in the high-price environment of 2022, only 29.3% of the energy purchased by the portfolio in Case C is sourced from the DAM, while the remaining 70.7% is still procured through the retail tariff. This balance also varies markedly across consumers, with the market share ranging

from 23.5% to 42.8% (Table A.1 in Appendix A). The bi-hourly consumer sources the largest fraction of its purchases from the market, providing an early indication of the consumer-level heterogeneity examined later in Section 4.5. Together with the limited arbitrage opportunity available to the 5 kWh storage units, this partial displacement of tariff procurement leaves the cumulative saving at 47.9% in Case C. This establishes the pre-reserve economic benchmark against which the substantially larger contribution introduced in Case D is assessed next.

4.3.3 BESS Scheduling and Reserve Provision

The step from Case C to Case D is the largest in the entire case study, reducing the annual variable cost from 4988 € to 622 €. This subsection examines the mechanism behind this reduction. Rather than lowering the cost of energy procurement, reserve provision introduces new portfolio-level revenues associated with reserve-band availability and expected activation.

Reserve provision is derived exclusively from the portfolio BESS units. Following the formulation of Section 3.5.2, each battery contributes upward and downward headroom around its scheduled operating point, and these individual contributions are aggregated into portfolio-level upward and downward reserve bands subject to the fixed two-to-one relationship between them. Non-zero reserve bands are scheduled in 94% of the hours of the year. This near-continuous participation is not imposed by the formulation but emerges from the optimisation, as retaining sufficient BESS headroom creates an additional source of revenue. Hours without reserve offers occur primarily when storage is already discharging close to its operational limit, leaving insufficient upward headroom to form the required reserve band.

Maintaining this reserve capability is not without an energy-side cost. Preserving the SOC margins required by the reserve bands changes the BESS schedule relative to Case C, increasing energy cycling and the associated round-trip losses. As Table 4.8 shows, the net energy cost consequently rises from 4988 € in Case C to 5201 € in Case D. The overall reduction in variable cost is instead driven by 4579 € of reserve revenue, of which roughly two thirds arise from band remuneration and one third from expected activation. This revenue is more than twenty times the additional energy-side cost incurred to enable reserve provision.

Table 4.8: Decomposition of the annual variable cost before and after reserve provision.

Component [€/year]	**Case C**	**Case D**
Energy purchase cost	6557	6778
Energy sale revenue	−1569	−1577
Net energy cost	**4988**	**5201** ▲
Reserve band revenue	-	−3003
Expected activation revenue	-	−1576
Total variable cost	**4988**	**622** ▼

Arrows show the change from Case C to Case D: ▲ higher, ▼ lower.

This decomposition makes the central economic mechanism explicit. The portfolio accepts a modest deterioration in its energy-side position in order to preserve reserve capability, but the associated band and expected activation revenues more than compensate for this additional cost. The dominant improvement from Case C to Case D therefore comes from monetising BESS availability rather than from reducing the cost of energy procurement.

These results indicate that, under the assumptions adopted in this case study, reserve participation becomes the primary economic driver of portfolio optimisation, with subsequent EMS layers acting mainly as refinements of the remaining energy cost.

4.3.4 Electric Vehicle Scheduling

The EV fleet represents the largest explicitly schedulable load in the portfolio, accounting for around 25,100 kWh of annual charging across the four consumers that own a vehicle. This energy requirement varies only marginally across configurations because each vehicle must satisfy its required departure SOC at the morning plug-out. Since V2G is disabled in this case study, the vehicles operate exclusively as flexible demand and never discharge to the grid. The main degree of freedom available to the EMS is therefore when this charging takes place within each vehicle's availability window.

Fig. 4.2 shows the resulting portfolio charging profile, averaged by hour of day over the full 2022 simulation, and illustrates how this timing changes as successive optimisation and market layers are introduced.

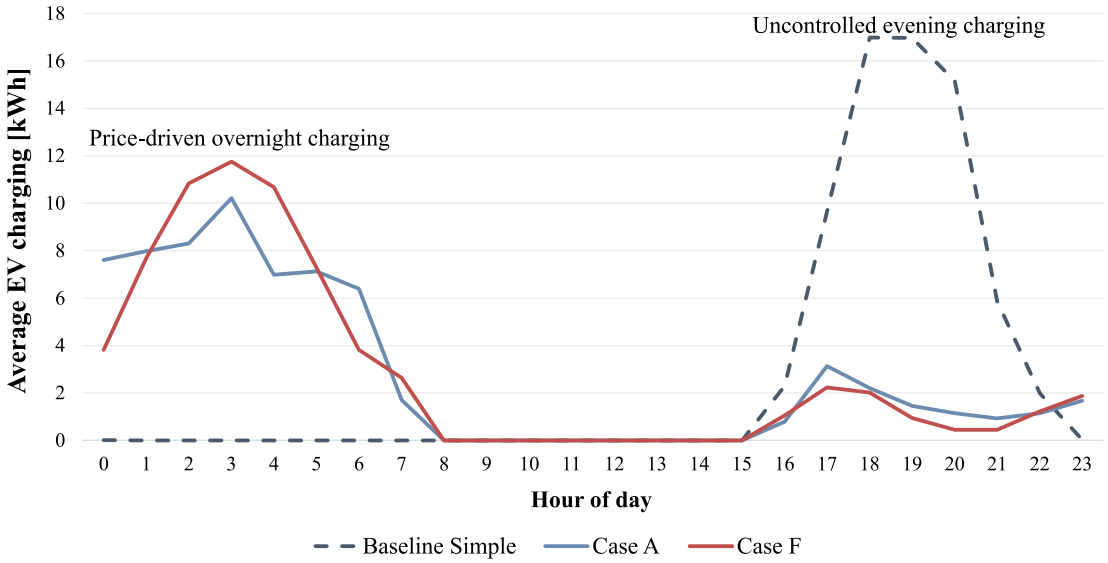


Figure 4.2: Average hourly portfolio EV charging profiles.

Under the passive baseline, charging is uncontrolled. Each vehicle draws power at its maximum rate as soon as it is plugged in, producing a pronounced evening peak shortly after the typical arrival time. The required energy is therefore delivered in a relatively short, high-power charging block without responding to price signals. Case A softens this evening peak rather than removing

it completely. Around 82% of the charging energy is shifted into the lower-cost overnight tariff periods, turning the concentrated evening block into a broader overnight plateau. The remaining 18% stays closer to plug-in, disproportionately associated with the consumer on the flat daily tariff, for whom the EMS has little price incentive to defer charging further into the night. This more distributed charging pattern increases the number of hours in which at least one EV is charging from 2,569 in the passive baseline to 4,287 in Case A.

Once market-based purchasing becomes available, the charging profile becomes more concentrated again. This transition begins in Case C, when the DAM is introduced on the buy side, and becomes more pronounced in Case F, where charging is concentrated in the deep overnight hours, peaking around 03:00, and the number of active charging hours falls to 3,256. Rather than responding only to the retail tariff structure, the EMS now schedules charging according to the lowest-cost opportunities available within the plug-in window under the broader procurement framework. Across all configurations, charging remains confined to the period between arrival in the late afternoon and the morning plug-out. The vehicles are therefore unavailable during the midday PV peak and cannot directly absorb local generation at that time, leaving the corresponding surplus to be managed through storage and grid exchanges. Despite the substantial redistribution of charging across the available hours, the required departure SOC is satisfied throughout the simulation.

The results therefore show that the rolling-horizon implementation preserves consumer mobility requirements while enabling economically driven charging decisions.

4.3.5 Demand-Side Flexibility Activation

DSF is introduced in Case E as a pure load-shifting mechanism. Following the formulation of Section 3.5.2, the EMS may raise or lower each consumer's demand within its data-driven flexibility bounds, but any increase must be matched by an equal decrease within the same day so that the total energy served is preserved. This daily neutrality, imposed by Equation 3.10, is maintained to numerical precision in the results. Over the year, the portfolio shifts approximately 3844 kWh upward and the same amount downward, corresponding to about 11% of its annual baseline consumption being redistributed in time with effectively zero net energy change. The extent of this reshaping is bounded by each consumer's flexibility envelope, which limits how far demand may deviate from its historical profile. Within these limits, flexibility is used in small amounts across almost the entire year rather than being concentrated in occasional large activation blocks. Since DSF receives no direct remuneration in the objective function, its economic value arises indirectly through its interaction with energy procurement, selling, storage, and reserve operation. Because neutrality is imposed separately for each day, the optimiser cannot advance or defer flexible demand across day boundaries. In this formulation, the temporal scope of DSF is therefore narrower than that of storage, whose SOC can carry stored energy across day boundaries. This restriction narrows the arbitrage opportunities available to flexible demand and is one factor limiting its economic contribution.

Fig. 4.3 shows how this reshaping is distributed over the day, averaged by hour across all days of 2022, with positive values representing net increases in demand and negative values representing net reductions.

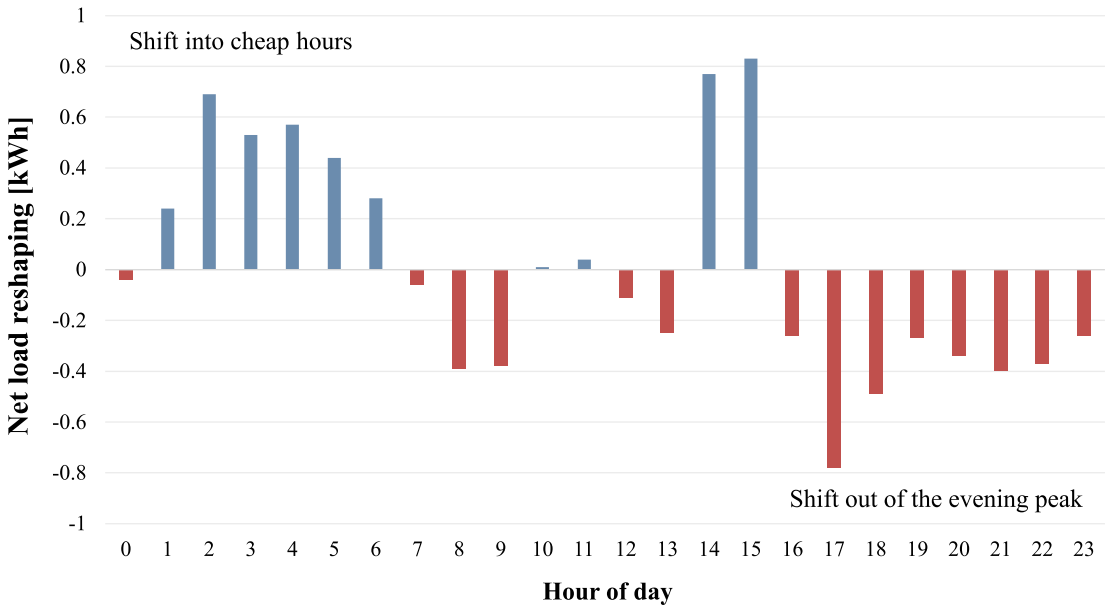


Figure 4.3: Average hourly portfolio net load reshaping in Case E.

Demand is shifted towards lower-cost periods, particularly during the early morning and the mid-afternoon window between 14:00 and 15:00, while consumption is reduced most strongly during the more expensive evening period from approximately 16:00 to 23:00. The early-morning and mid-afternoon increases reflect favourable procurement opportunities, whereas the evening reductions occur when energy has greater value in avoiding imports or supporting exports. The resulting profile therefore combines valley-filling and peak-shaving while reflecting the interaction between hourly procurement conditions and the opportunity value of energy within the portfolio.

The resulting saving is modest in absolute terms, lowering the variable cost from 622 € in Case D to 450 € in Case E, but the mechanism behind it is instructive. Although demand is shifted towards lower-cost periods, the improvement does not arise from a reduction in total procurement expenditure. Annual grid purchases increase and the associated purchase cost rises by about 68 €. The larger benefit comes from the selling side. By moving consumption away from hours when locally generated or stored energy can instead be exported, flexibility releases roughly 1150 kWh of additional energy for sale, increasing annual sales revenue by approximately 233 €. Together with the small change in reserve revenue, these opposing effects produce the 172 € reduction in variable cost. DSF therefore creates value not by reducing the portfolio's energy requirement, but by redistributing demand so that energy is consumed, retained, or sold at more economically favourable times.

These results demonstrate that the consumer-specific flexibility envelopes can be consistently integrated into the EMS while preserving daily energy neutrality.

4.3.6 Bilateral Procurement

The final layer, introduced in Case F, gives the aggregator access to a bilateral supply contract at a fixed price of 0.1369 €/kWh for all consumers. Unlike the DAM and the time-differentiated retail tariffs, the bilateral price remains constant throughout the year. The contract therefore provides a third procurement option whose value depends on how its fixed price compares with the available alternatives and on the contractual limits governing its use.

The contract supplies 4482 kWh over the year, corresponding to about 7.5% of the energy purchased by the portfolio, and is used in 3124 hours. At 0.1369 €/kWh, it is cheaper than the 2022 average DAM price but more expensive than the lowest tariff bands, such as the 0.1072 €/kWh tri-hourly off-peak rate reported in Section 4.1.1. Its use is concentrated in two main windows where the fixed bilateral price is particularly competitive with the available alternatives. The evening period from 18:00 to 23:00 accounts for 57% of the contracted volume, while the morning period from 07:00 to 11:00 accounts for a further 26%. The remaining volume is distributed across other hours, partly reflecting the minimum daily procurement requirement. Since the purchase mix changes only marginally across Cases C to E, Case C is used in Table 4.9 as a representative pre-bilateral reference. The introduction of bilateral procurement in Case F then changes the portfolio mix to 68.4% tariff, 24.1% market, and 7.5% bilateral purchases.

Table 4.9: Portfolio energy purchase shares by procurement channel in Cases C and F.

Buy channel	**Case C**	**Case F**
Tariff	70.7%	68.4%
Market	29.3%	24.1%
Bilateral	-	7.5%

Following the formulation of Section 3.5.2, the bilateral contract is also subject to daily procurement obligations, including a minimum volume that must be purchased and a maximum that cannot be exceeded. The minimum and maximum limits bind on 87 and 138 days, respectively. The minimum can force purchases when another channel is marginally cheaper, whereas the maximum caps otherwise attractive bilateral procurement. The aggregate share also masks substantial heterogeneity across consumers. Bilateral use remains below 1% for the two flat-tariff consumers but exceeds 13% for the bi-hourly consumer, as shown in Table A.1 in Appendix A. The overall economic contribution is nevertheless the smallest of the successive EMS layers, reducing the annual variable cost by 166 €, from 450 € in Case E to 284 € in Case F. This improvement arises mainly from the procurement side, as contracted energy replaces more expensive purchases whenever the bilateral option is economically favourable and the contractual limits permit its use.

The modest saving observed in Case F reflects the deterministic value of the fixed-price contract under the observed 2022 prices and contractual limits. Its potential hedging value under price uncertainty is not quantified by the perfect-foresight formulation and is revisited in Section 4.6.

4.4 Seasonal Analysis

The results discussed so far are annual aggregates, which show the overall effect of the EMS configurations but do not reveal how the portfolio operates under contrasting conditions throughout the year. Since the full-year simulation is implemented through consecutive daily rolling-horizon cycles, it is relevant to examine whether the same formulation produces coherent operating decisions under markedly different seasonal conditions. For this purpose, two representative weekdays are analysed under the fully active Case F, selected based on their proximity to the median demand and PV output of their respective months: 5 July for summer and 18 January for winter.

Fig. 4.4 shows the hourly portfolio energy balance on both days, with grid imports and BESS discharge above the axis and grid exports, BESS charging, and EV charging below it.

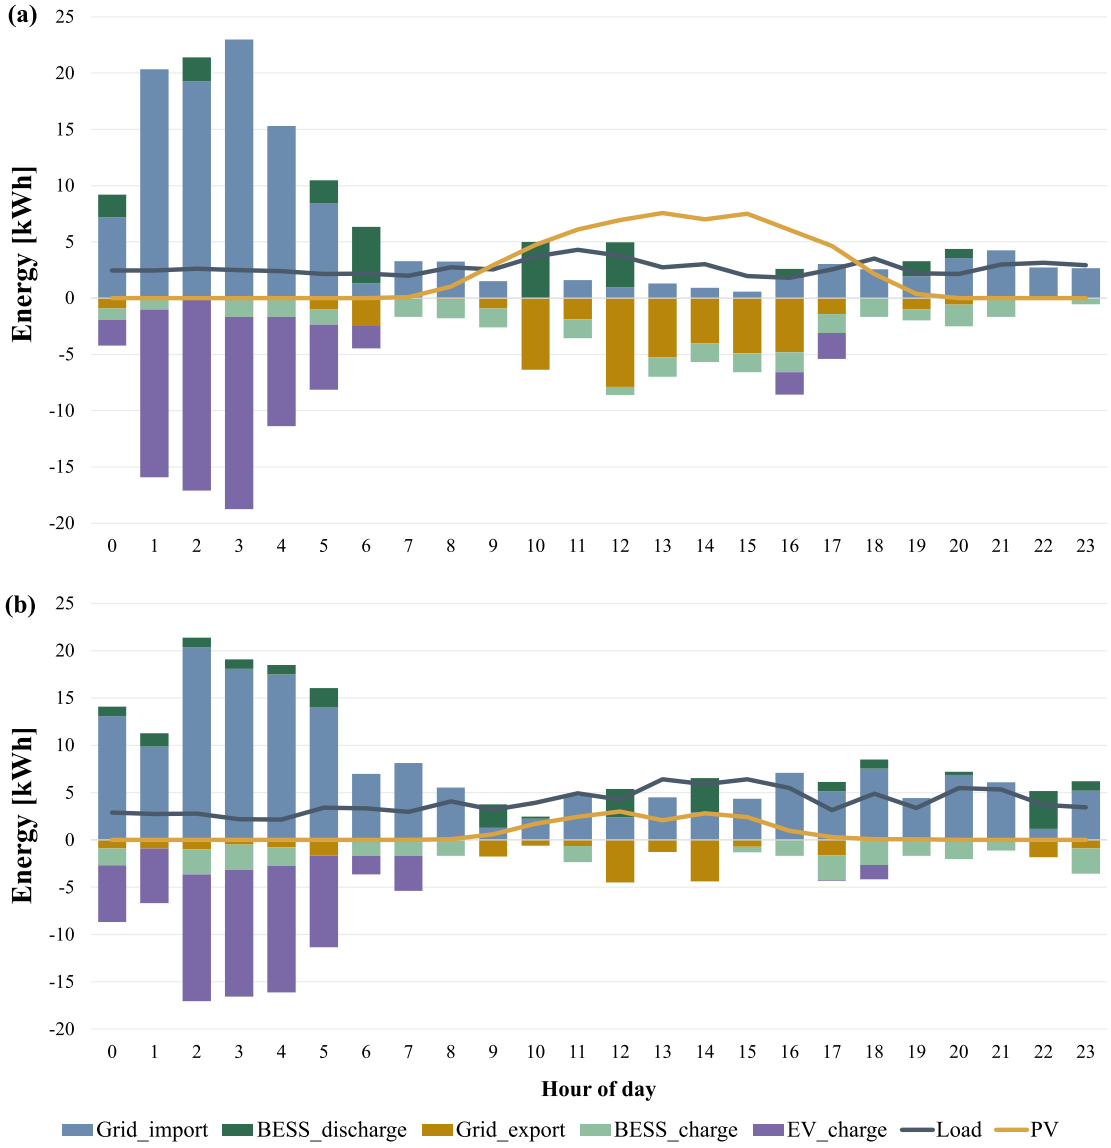


Figure 4.4: Daily portfolio operation under Case F: (a) summer day and (b) winter day.

The comparison first highlights the operating patterns that remain relatively stable across the two seasonal conditions. EV charging is concentrated in a similar overnight window on both days, with more than 90% of the daily charging energy delivered between midnight and 07:00, namely 94% in summer and 92% in winter. This overnight charging block drives the early-morning import peaks, which reach 23.0 kWh on the summer day and 20.4 kWh on the winter day. Reserve availability is similarly stable. The daily sum of the hourly reserve bands is around 126 kWh upward and 63 kWh downward on both days, differing by less than 1%. This suggests that, for the two representative days, the available reserve headroom is governed predominantly by the technical capability and operating margins of the BESS fleet and is less sensitive to seasonal PV conditions than the portfolio's energy exchanges.

The strongest seasonal contrast appears in the interaction between local generation, storage, and grid procurement. On the summer day, PV generation totals 53 kWh between 09:00 and 17:00, creating sufficient surplus for 38 kWh to be exported over the same period. On the winter day, generation during these hours falls to around 30% of the summer level, substantially reducing the available surplus and increasing reliance on grid imports around the middle of the day. The BESS schedule adapts accordingly, balancing energy arbitrage with the need to preserve reserve headroom. These different operating conditions also reshape the procurement mix. On the summer day, the DAM supplies approximately 35.7% of purchased energy, while the retail tariff and bilateral contract account for around 58.6% and 5.7%, respectively. On the winter day, the DAM is not used for procurement, with approximately 89% of purchases supplied through the retail tariff and the bilateral contract covering the remainder.

A further seasonal difference appears in the remuneration of reserve provision. Although the quantity of reserve offered remains almost unchanged, the composition of the associated revenue differs substantially. The average band price on the summer day is nearly five times the winter value, whereas activation prices and expected activation ratios are higher on the winter day. These effects partially offset each other, resulting in reserve revenues of a comparable order on both days, at 13.5 € in summer and 10.8 € in winter, but generated through a different mix. Band remuneration dominates on the summer day, while the expected activation component becomes more important in winter.

These mechanisms ultimately translate into different daily economic outcomes. The portfolio records a slightly negative variable cost on the summer day, at -4.4 €, meaning that selling and reserve revenues exceed the corresponding procurement cost. The winter day instead results in a positive variable cost of 6.5 €. The difference of approximately 11 € is mainly associated with higher procurement expenditure, followed by lower reserve income and lower selling revenue on the winter day. DSF use remains of a similar magnitude, with 11.1 kWh shifted in each direction on the summer day and 9.8 kWh in each direction on the winter day.

Taken together, the two representative days illustrate how the same EMS formulation adapts to contrasting seasonal conditions while preserving its main operational requirements. EV charging remains concentrated within the available overnight window, reserve availability is broadly stable, and the magnitude of DSF use changes only moderately. In contrast, PV surplus, BESS energy

scheduling, grid imports and exports, the procurement mix, and the composition of reserve revenues respond more strongly to the prevailing conditions. These representative results therefore show that the rolling-horizon formulation produces coherent operating decisions under the two contrasting seasonal conditions without requiring seasonal retuning.

4.5 Consumer-Level Analysis

The seasonal analysis still treats the five consumers as a single portfolio, although they differ substantially in tariff structure, EV ownership and characteristics, demand, and PV generation while sharing identical BESS specifications. This section therefore revisits the same two questions from an individual-consumer perspective: how the value created through coordinated operation and energy-market access varies across the portfolio, and how the additional value introduced by reserve participation, DSF, and bilateral procurement is distributed among consumers with different energy needs and cost structures.

Fig. 4.5 compares the annual variable cost of each consumer under Baseline Simple and Case F, while the complete progression across all eight configurations is reported in Table A.2 in Appendix A. Under Case F, Users 1, 2, and 5 reach negative variable costs because their attributed selling and reserve revenues exceed their annual procurement costs under the adopted reporting allocation. These negative values should be interpreted as negative net variable operating costs within the modelled scope rather than as evidence of overall annual profitability. The comparison excludes the fixed contracted-power charge and does not include asset investment or other implementation costs.

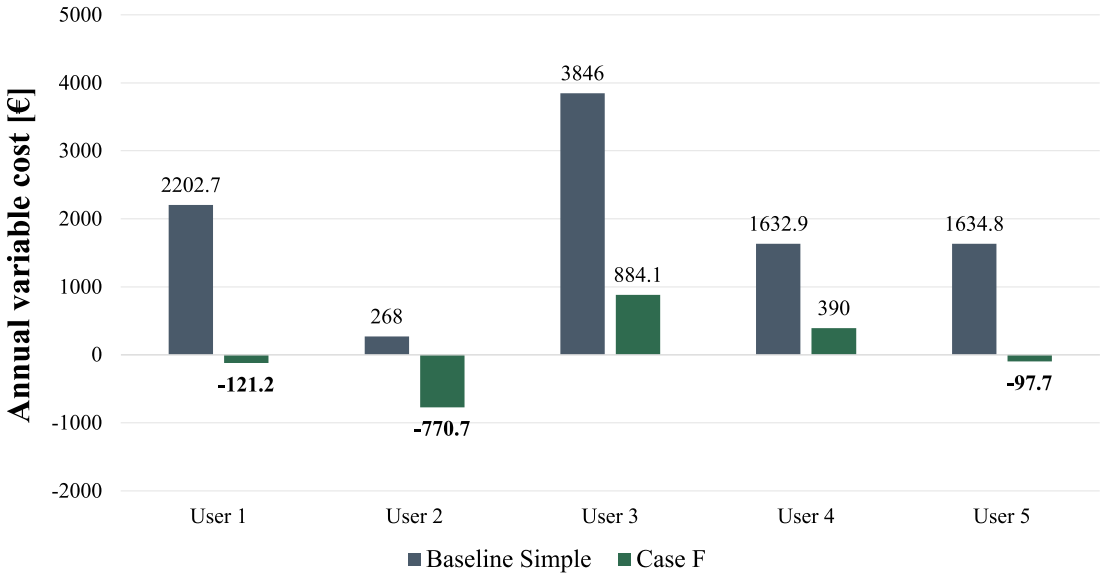


Figure 4.5: Annual variable cost per consumer under Baseline Simple and Case F.

For the first question, the value of coordinated asset scheduling differs markedly across consumers. Case A isolates the effect of optimised BESS scheduling and EV smart charging, where

applicable, under retail-tariff procurement before remunerated selling or market-based purchasing are introduced. User 1, on the tri-hourly tariff, reduces its variable cost by 50.6% relative to Baseline Simple, whereas User 4, on the flat daily tariff, achieves a reduction of only 5.1%. Both consumers have the same BESS specification, yet the time-varying tariff available to User 1 creates substantially greater opportunities for shifting storage and EV charging towards economically favourable periods. By contrast, the flat tariff removes time-of-use arbitrage within the retail price itself, leaving fewer opportunities for coordinated scheduling to create value. The comparison therefore shows that the economic benefit of identical flexible assets depends strongly on the price structure and operating conditions faced by each consumer.

The second question concerns how the additional value created at aggregator level is distributed among consumers. The dominant mechanism is again reserve provision, identified in Section 4.3.3 as the main economic driver of the portfolio results. For consumer-level accounting, the aggregated reserve band revenue is allocated ex post according to each BESS's contribution to the offered reserve band. This allocation affects only the attribution of financial results and does not alter the optimised dispatch or the aggregate portfolio outcome. Together with the corresponding expected activation revenues, total reserve income remains relatively similar across the five consumers, ranging from 904 € to 933 € per year, a spread of less than 3.3%. This similarity reflects the identical technical specifications of the five BESS units, although small differences remain because their available reserve contributions depend on their individual operating schedules.

The significance of this similar reserve income nevertheless depends on each consumer's underlying cost. User 2 begins with a Baseline Simple cost of only 268 € and saves just 4.0% in Case A, but reaches a negative cost in Case F, with attributed revenues exceeding its variable costs by approximately 771 €. User 3 begins with the highest cost, at 3846 €, and retains a positive cost of 884 € in Case F despite a saving of 77.0%. The decomposition between band and expected activation revenues is reported in Table A.3 in Appendix A. Similar reserve contributions can therefore have very different economic significance depending on the underlying cost structure.

A further source of heterogeneity is DSF. User 4 shifts 15.1% of its annual load in Case E, the largest relative share in the portfolio, whereas User 3 shifts only 5.7% despite having the highest demand. These differences reflect the consumer-specific flexibility envelopes and their interaction with EV charging, storage operation, and the remaining technical constraints. Available flexibility therefore depends not simply on demand size, but on how much demand can be shifted within each consumer's operating profile.

Beyond cost, optimisation and market participation are also assessed through local energy use using the Self-Consumption Ratio (SCR) and the Self-Sufficiency Ratio (SSR). Since the model does not track the origin of energy discharged from the BESS, the self-consumed PV energy of consumer i in period t is approximated for reporting as contemporaneous local generation not exported, $E^{SC}_{i,t} = \max\left(0,\ \widehat{PV}_{i,t} - E^{sell}_{i,t}\right)$:

$$\mathrm{SCR}_i = \frac{\sum_t E^{SC}_{i,t}}{\sum_t \widehat{PV}_{i,t}} \times 100, \qquad \mathrm{SSR}_i = \frac{\sum_t E^{SC}_{i,t}}{\sum_t \left(\widehat{L}_{i,t} + E^{EV,ch}_{i,t}\right)} \times 100. \tag{4.2}$$

The resulting SCR and SSR values should therefore be interpreted as reporting proxies rather than as exact provenance-based indicators. Table 4.10 reports both indicators for Baseline Simple and Case F, matching the endpoint comparison used in Fig. 4.5.

Table 4.10: SSR and SCR per consumer under Baseline Simple and Case F.

	SSR		SCR	
Consumer	**Baseline Simple**	**Case F**	**Baseline Simple**	**Case F**
User 1	14.0%	13.6%	48.1%	46.6%
User 2	26.5%	23.0%	88.0%	76.5%
User 3	3.6%	3.3%	89.9%	82.4%
User 4	16.9%	15.9%	64.7%	60.9%
User 5	15.9%	16.4%	52.1%	53.6%

Viewed only at these two endpoints, the changes appear relatively modest and are not uniform across consumers. User 5, for example, ends Case F with both indicators slightly above its Baseline Simple values, whereas the remaining consumers show moderate reductions. The full progression across all eight configurations, reported in Tables A.4 and A.5 in Appendix A, reveals a much clearer underlying pattern. Activating the BESS increases local use of PV substantially relative to Baseline Simple, with SCR and SSR generally reaching their highest values around Baseline BESS or Case A, when storage is operated primarily to support local demand and tariff optimisation. Once remunerated surplus selling is introduced in Case B, both indicators decrease for every consumer because exporting begins to compete with local self-consumption for the economic value of available generation and stored energy.

This reversal is particularly pronounced for User 1, whose SCR falls from 82.1% in Case A to 51.1% in Case B. For User 1, the strong incentive to use local flexibility for tariff optimisation in Case A is therefore partly displaced once remunerated export becomes available and selling local energy acquires an explicit economic value. The corresponding reduction is smallest for User 3, whose comparatively limited PV generation restricts the scope for reallocating locally generated energy. Subsequent configurations modify these indicators further, but they do not restore the high self-consumption levels reached before remunerated selling was introduced. The comparison between Baseline Simple and Case F alone therefore conceals an important intermediate effect: optimisation initially increases local self-consumption, after which market participation changes the economic priority from maximising local use towards maximising portfolio value.

Taken together, the consumer-level results show that the benefits of aggregation are positive but uneven in relative terms. Identical BESS assets generate broadly similar reserve revenues, yet the importance of that income differs substantially because consumers begin with different tariffs, demand levels, mobility requirements, and local generation profiles. At the same time, maximising portfolio-level economic value does not necessarily maximise individual self-consumption or self-sufficiency, particularly once remunerated exports become available. This uneven distribution of economic and energetic outcomes motivates the trade-off discussion in the following section.

4.6 Trade-Off Discussion

The preceding sections have shown that the aggregated EMS can substantially reduce the portfolio's variable operating cost, with reserve provision emerging as the dominant source of additional value and with the benefits varying across seasons and consumers. Taken together, however, these results also reveal five trade-offs that qualify how directly the case-study outcomes can be transferred to practical deployment. Table 4.11 summarises the principal benefit and limitation associated with each progressive configuration from Case A to Case F, before the main cross-cutting trade-offs are discussed in greater detail.

Table 4.11: Principal benefit and limitation of each progressive EMS case.

Case	Principal benefit	Principal limitation
Case A	Coordinated BESS and EV scheduling substantially reduces variable operating cost under tariff-based procurement.	The achievable benefit depends strongly on the consumer's tariff and operating profile, while investment costs are outside the case-study objective.
Case B	Remunerated exports create an additional revenue stream from surplus and stored energy.	Selling competes with local energy use, reducing SCR and SSR relative to Case A.
Case C	DAM procurement broadens the available purchasing options and further reduces operating cost.	Its value depends on market-price opportunities, perfect foresight within the optimisation window, and the limited arbitrage capability of the installed storage.
Case D	Reserve provision delivers by far the largest incremental economic benefit.	The result depends on relaxed participation conditions and on expected activation rather than physically simulated reserve delivery.
Case E	Data-driven DSF adds operational flexibility while preserving daily energy neutrality.	Its additional economic contribution is modest and depends on estimated consumer-specific flexibility envelopes.
Case F	Bilateral procurement introduces an additional fixed-price purchasing option.	Contractual volume limits restrict its use, while the deterministic formulation does not capture its potential hedging value under uncertainty.

The first and most consequential trade-off concerns the large value obtained from reserve provision and the conditions required to access it. Case D produces the dominant improvement in the entire case study, but the five-consumer portfolio is far below the scale required for standalone participation in the mFRR market under the market conditions considered in this dissertation. The relaxed participation threshold adopted in the model should therefore be interpreted as a proxy for participation within a wider aggregation rather than as direct market eligibility of the portfolio itself. The reserve revenues consequently represent the economic potential of distributed BESS capacity when access through a sufficiently large and qualified aggregation is available.

This result is also conditioned by how reserve activation is represented. Revenue is based on expected activation derived from historical activation prices and ratios, while physical activation, delivery uncertainty, imbalance exposure, and potential non-delivery penalties are not simulated.

Reserve provision remains the most valuable layer under the assumptions of the case study, but the magnitude of that benefit depends strongly on market access and on the simplified representation of activation. It should therefore not be interpreted as revenue that this five-consumer portfolio could independently reproduce under real operating conditions.

A second trade-off concerns the deterministic, perfect-foresight formulation. Each 32-hour rolling-horizon optimisation is solved using realised demand, generation, and market data as known inputs. This allows the EMS to schedule storage and EV charging around known price opportunities without accounting for forecast errors, creating a more favourable information environment than would exist in real operation. At the same time, this assumption removes any explicit value associated with protection against future price uncertainty. This is particularly relevant to the bilateral contract, whose contribution reflects only its deterministic price advantage over alternative procurement channels. Its potential value as a hedge against future market volatility cannot be captured when future prices are already known. The reported balance between market-based and contractual procurement should therefore be interpreted in the context of this deterministic setting rather than as a complete comparison of their real-world value.

A third trade-off, revealed at consumer level, concerns the value created by aggregation and how evenly that value is shared. Reserve revenues are broadly similar across the five consumers because they have identical BESS specifications and the aggregated band revenue is allocated ex post according to each unit's contribution to the reserve band. Their economic significance, however, differs substantially because the consumers begin from very different cost levels. A similar reserve contribution can therefore transform a small consumer's annual position into a net gain while leaving a larger consumer with a positive, although substantially reduced, cost. Other mechanisms are more strongly shaped by individual conditions, as illustrated by the substantial differences in bilateral procurement across tariff structures. A practical aggregator would therefore face a distributional question that portfolio-level optimisation alone cannot resolve, namely how the value created collectively should be allocated in a way that reflects contribution while remaining attractive to consumers with different energy positions and potential gains.

A fourth trade-off concerns economic optimisation and household energy autonomy. When the BESS is used primarily to support local demand and tariff optimisation, self-consumption and self-sufficiency generally increase. Once remunerated selling becomes available, part of the role of storage shifts towards monetising energy through external markets, and both indicators decline relative to Case A for every consumer. The same market access that creates additional financial value can therefore reduce the reported self-consumption and self-sufficiency of individual households. This is not a flaw in the formulation, since the EMS minimises variable operating cost and was not designed to maximise self-sufficiency. It does, however, show that the financial case for aggregation and the objective of greater household energy autonomy can point in different directions, particularly for consumers who value local energy use in addition to monetary savings.

The fifth trade-off is between the additional value created by successive EMS layers and the complexity required to implement them. Reserve provision captures most of the incremental economic value obtained after full energy-market access, while DSF and bilateral procurement pro-

vide smaller subsequent improvements. This does not make the later layers redundant, since each contributes through a distinct mechanism, but their measured marginal value depends on the sequence adopted in the case study. By the time Cases E and F are introduced, reserve provision has already captured most of the observed incremental benefit under the adopted sequence. Their smaller contributions should therefore not be interpreted as evidence that these mechanisms would always have limited value under different market conditions or in a different portfolio. Moreover, the case study compares variable operating costs and does not quantify asset investment or the additional forecasting, communication, contractual, and operational costs associated with implementation. In practice, the marginal benefit of each additional layer would need to justify these requirements.

None of these trade-offs undermine the central finding of this dissertation, that an aggregated, multi-market EMS can unlock value that individual consumers and simpler operating strategies cannot reach on their own. They do, however, mark the boundary between what the model demonstrates and what a real deployment would still need to resolve. These issues include access to balancing markets, operation under uncertainty, credible reserve activation and settlement, a fair allocation of portfolio-level value, the balance between financial return and household energy autonomy, and whether the marginal benefit of additional EMS layers justifies their implementation requirements. Chapter 5 returns to these limitations when assessing the main findings of the dissertation and the corresponding directions for future work.

Chapter 5

Conclusion

This chapter closes the dissertation with a critical assessment of the aggregated EMS, the case study results, and the trade-offs they expose. It summarises the main findings, states the contributions of this work, and identifies the limitations that motivate the directions for future work proposed at the end of the chapter.

5.1 Conclusions

This dissertation developed and extended an aggregated EMS for B2B and B2C consumer portfolios and evaluated it through a residential B2C case study, mapping the trade-offs that emerge as distributed resources are progressively coordinated across multiple energy, reserve, and procurement mechanisms.

The proposed EMS extends a pre-existing single-site optimisation tool into a multi-consumer deterministic MILP formulation that coordinates multiple consumers over a rolling horizon, preserving each consumer's individual physical feasibility while coordinating selected decisions at portfolio level. Intertemporal consistency across rolling cycles is maintained through BESS SOC carry-over and dedicated treatment of ongoing overnight EV sessions. In parallel, a data-driven DSF module estimates consumer-specific flexibility bounds directly from historical consumption, without assuming a fixed or arbitrary flexibility margin.

The case study, built on real 2022 consumption and PV generation data from five residential consumers, evaluated two baselines and six progressively enriched EMS cases. Coordinated scheduling of storage and electric mobility, combined with remunerated selling and market-based procurement, reduced the portfolio's annual variable operating cost by 47.9% relative to the fully passive baseline. mFRR reserve provision proved to be the single most valuable additional layer, raising the cumulative saving to 93.5%. This improvement arose primarily from remunerating available BESS capacity, complemented by expected activation revenues, rather than from reducing the underlying energy procurement cost. At the scale considered here, this reserve value should therefore be interpreted as the economic potential available through a sufficiently large

and qualified aggregation under the adopted expected-activation assumptions, rather than as revenue directly realisable by the five-consumer portfolio in isolation. DSF and bilateral procurement provided further, more moderate improvements, bringing the cumulative saving to 97.0% in the fully active Case F. The two representative seasonal days also showed coherent EMS operation under contrasting conditions without requiring seasonal retuning. These results demonstrate that the economic value of aggregation extends well beyond energy arbitrage, increasingly relying on coordinated participation across complementary market mechanisms.

Examined at consumer level, the results show that the economic value created by the EMS is not distributed uniformly across the portfolio. Tariff structure, demand, electric mobility, local generation, and available flexibility all influence how much each consumer benefits from coordinated scheduling and market access. Even with identical BESS specifications, the value obtained from optimisation differs substantially across consumers. Reserve revenues are comparatively similar in absolute terms because the batteries are technically identical and the aggregated band revenue is attributed ex post according to individual reserve contribution, yet their relative economic impact differs markedly with each consumer's initial cost position. The analysis further shows that financial optimisation and household energy autonomy are not necessarily aligned, since remunerated exporting can reduce the reported self-consumption and self-sufficiency indicators while lowering operating cost. These findings, together with the deterministic nature of the formulation, the regulatory and modelling assumptions underlying reserve participation, and the diminishing marginal value observed for successive EMS layers under the adopted case sequence, were consolidated into five trade-offs that qualify how far the case-study results can be generalised beyond this specific portfolio and market context.

The central contribution of this dissertation is the aggregated MILP formulation itself. The resulting EMS operates over a rolling horizon and coordinates storage, electric mobility, data-driven flexibility, energy-market transactions, reserve provision, and bilateral procurement across a multi-consumer portfolio while preserving each consumer's technical constraints. By embedding these decisions within a common optimisation problem, the framework represents aggregation as coordinated operation rather than as a simple ex post combination of independently optimised consumer schedules. A further contribution lies in the progressive simulation methodology adopted throughout the case study, in which successive EMS layers are activated to quantify their incremental effects within the adopted sequence and systematically identify the operational and economic trade-offs associated with multi-market participation.

An earlier, non-aggregated version of the underlying formulation was presented at an international conference, while the formulation developed in this dissertation extends that work to coordinated multi-consumer operation and the additional market mechanisms evaluated here.

Taken together, these results show that an aggregated, multi-market EMS can unlock value that individual consumers and simpler heuristics cannot reach on their own, while making clear that realising this value in practice depends on regulatory access, a fair allocation of portfolio-level benefits, and an explicit treatment of the uncertainty left unaddressed by the deterministic formulation.

5.2 Future Work

Several directions emerge directly from the trade-offs and modelling choices discussed across this dissertation, namely:

- Extending the case study to B2B portfolios, broadening the range of asset configurations and consumption profiles evaluated by the framework;
- Validating reserve participation within a sufficiently large and qualified aggregation capable of meeting the applicable mFRR participation requirements, while explicitly representing physical activation, delivery performance, and potential non-delivery consequences;
- Reformulating the model under stochastic or robust optimisation to represent uncertainty in demand, PV generation, market prices, and reserve activation, while explicitly assessing the hedging value of bilateral procurement that the current deterministic formulation cannot capture;
- Testing more conservative and scenario-based representations of expected mFRR activation ratios, complementing the expected-value representation adopted in the present case study and assessing the sensitivity of reserve value to alternative activation conditions;
- Exploring alternative reserve-revenue allocation rules by comparing physical-contribution-based attribution with approaches reflecting each consumer's energy position, individual benefit, or broader participation incentives;
- Exploring objective formulations beyond portfolio operating-cost minimisation, including multi-objective and social-welfare-oriented approaches that consider self-sufficiency, consumer utility and comfort, benefit distribution, and broader system-level impacts of aggregator decisions;
- Enabling and evaluating V2G operation, disabled in the current case study, and assessing its effect on energy scheduling, reserve availability, consumer mobility requirements, and self-consumption;
- Increasing the temporal resolution of the rolling-horizon simulation to 15 minutes, aligned with the 15-minute market time unit used in European day-ahead market coupling and allowing faster operational dynamics to be represented more accurately;
- Extending the economic assessment beyond variable operating cost to include asset investment and the additional forecasting, communication, metering, market-access, contractual, and operational costs associated with implementation, allowing the economic benefits of additional EMS layers to be assessed against their full deployment requirements.

Together, these research directions provide a natural path for extending the proposed EMS towards more realistic, scalable, and market-ready aggregation frameworks capable of supporting the evolving role of distributed flexibility in future electricity systems.

References

[1] International Energy Agency, "Net zero roadmap: A global pathway to keep the 1.5 °C goal in reach," International Energy Agency, Tech. Rep., 2023, 2023 Update. [Online]. Available: https://www.iea.org/reports/net-zero-roadmap-a-global-pathway-to-keep-the-15-c-goal-in-reach.

[2] European Commission, "The European Green Deal," European Commission, Tech. Rep. COM(2019) 640 final, 2019. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52019DC0640.

[3] European Commission, "Fit for 55: Delivering the EU's 2030 climate target on the way to climate neutrality," European Commission, Tech. Rep. COM(2021) 550 final, 2021. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021DC0550.

[4] B. Petrovich, "European electricity review 2026," Ember, Report, Jan. 22, 2026. Accessed: Jul. 29, 2026. [Online]. Available: https://ember-energy.org/latest-insights/european-electricity-review-2026/.

[5] Ember. "Ember electricity data explorer," Accessed: Jul. 29, 2026. [Online]. Available: https://ember-energy.org/data/electricity-data-explorer/.

[6] F. Plaum, R. Ahmadiahangar, A. Rosin, and J. Kilter, "Aggregated demand-side energy flexibility: A comprehensive review on characterization, forecasting and market prospects," *Energy Reports*, vol. 8, pp. 9344–9362, 2022. DOI: 10.1016/j.egyr.2022.07.038.

[7] European Commission, "REPowerEU plan," European Commission, Tech. Rep. COM(2022) 230 final, 2022. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52022DC0230.

[8] C. Le Coq, A. R. Bennato, D. Duma, and E. Lazarczyk, "Flexibility in the energy sector," Centre on Regulation in Europe, Tech. Rep., 2025. [Online]. Available: https://cerre.eu/publications/flexibility-in-the-energy-sector/.

[9] M. M. Shepherd and S. Mohagheghi, "Demand flexibility: A review of quantification methods, models, and required data," *Applied Sciences*, vol. 14, no. 23, p. 11 276, 2024. DOI: 10.3390/app142311276.

[10] P. Crespo del Granado et al., "Flexibility characterization, aggregation, and market design trends with a high share of renewables: A review," *Current Sustainable/Renewable Energy Reports*, vol. 10, pp. 12–21, 2023. DOI: 10.1007/s40518-022-00205-y.

[11] P. Hernández Martínez, "Lithium-ion battery energy storage systems: Global market assessment and business model development in spain," Bachelor's thesis, Universidad Politécnica de Madrid, Escuela Técnica Superior de Ingenieros de Minas y Energía, 2020. [Online]. Available: https://oa.upm.es/63732/.

[12] J. P. Chaves, M. Troncia, C. Damas Silva, and G. Willeghems, "Overview of market designs for the procurement of system services by DSOs and TSOs," OneNet Project, Horizon 2020, Deliverable D3.1, 2021. [Online]. Available: https://onenet-project.eu/wp-content/uploads/2022/10/D31-Overview-of-market-designs-for-the-procurement-of-system-services-by-DSOs-and-TSOs.pdf.

[13] N. Mišljenović, M. Žnidarec, G. Knežević, D. Šljivac, and A. Šumper, "A review of energy management systems and organizational structures of prosumers," *Energies*, vol. 16, no. 7, p. 3179, 2023. DOI: 10.3390/en16073179.

[14] S. Mischos, E. Dalagdi, and D. Vrakas, "Intelligent energy management systems: A review," *Artificial Intelligence Review*, vol. 56, pp. 11 635–11 674, 2023. DOI: 10.1007/s10462-023-10441-3.

[15] K. Poplavskaya and L. de Vries, "Aggregators today and tomorrow: From intermediaries to local orchestrators?" In *Behind and Beyond the Meter: Digitalization, Aggregation, Optimization, Monetization*, F. Sioshansi, Ed., Elsevier, 2020, pp. 105–135. DOI: 10.1016/B978-0-12-819951-0.00005-0.

[16] T. Schittekatte, V. Deschamps, and L. Meeus, "The regulatory framework for independent aggregators," *The Electricity Journal*, vol. 34, no. 6, p. 106 971, 2021. DOI: 10.1016/j.tej.2021.106971.

[17] M. Adham, S. Keene, and R. B. Bass, "Distributed energy resources: A systematic literature review," *Energy Reports*, vol. 13, pp. 1980–1999, 2025. DOI: 10.1016/j.egyr.2025.01.026.

[18] H. M. Rouzbahani, H. Karimipour, and L. Lei, "A review on virtual power plant for energy management," *Sustainable Energy Technologies and Assessments*, vol. 47, p. 101 370, 2021. DOI: 10.1016/j.seta.2021.101370.

[19] H. Gao, T. Jin, C. Feng, C. Li, Q. Chen, and C. Kang, "Review of virtual power plant operations: Resource coordination and multidimensional interaction," *Applied Energy*, vol. 357, p. 122 284, 2024. DOI: 10.1016/j.apenergy.2023.122284.

[20] M. Kaiss, Y. Wan, D. Gebbran, C. U. Vila, and T. Dragičević, "Review on virtual power plants/virtual aggregators: Concepts, applications, prospects and operation strategies," *Renewable and Sustainable Energy Reviews*, vol. 211, p. 115 242, 2025. DOI: 10.1016/j.rser.2024.115242.

[21] J. Leitão, P. Gil, B. Ribeiro, and A. Cardoso, "A survey on home energy management," *IEEE Access*, vol. 8, pp. 5699–5722, 2020. DOI: 10.1109/ACCESS.2019.2963502.

[22] D. Mariano-Hernández, L. Hernández-Callejo, A. Zorita-Lamadrid, O. Duque-Pérez, and F. Santos García, "A review of strategies for building energy management system: Model predictive control, demand side management, optimization, and fault detect & diagnosis," *Journal of Building Engineering*, vol. 33, p. 101 692, 2021. DOI: 10.1016/j.jobe.2020.101692.

[23] T. H. B. Huy, H. T. Dinh, and D. Kim, "Multi-objective framework for a home energy management system with the integration of solar energy and an electric vehicle using an augmented ε-constraint method and lexicographic optimization," *Sustainable Cities and Society*, vol. 88, p. 104 289, 2023. DOI: 10.1016/j.scs.2022.104289.

[24] M. Tostado-Véliz, S. Mouassa, and F. Jurado, "A MILP framework for electricity tariff-choosing decision process in smart homes considering 'happy hours' tariffs," *International Journal of Electrical Power & Energy Systems*, vol. 131, p. 107 139, 2021. DOI: 10.1016/j.ijepes.2021.107139.

[25] M. A. A. Abdalla, M. Wang, B. Wang, A. M. Ishag, and B. Saleh, "Double-layer home energy management strategy for increasing PV self-consumption and cost reduction through appliances scheduling, EV, and storage," *Energy Reports*, vol. 10, pp. 3494–3518, 2023. DOI: 10.1016/j.egyr.2023.10.019.

[26] A. C. Duman, H. S. Erden, Ö. Gönül, and Ö. Güler, "Optimal sizing of PV-BESS units for home energy management system-equipped households considering day-ahead load scheduling for demand response and self-consumption," *Energy and Buildings*, vol. 267, p. 112 164, 2022. DOI: 10.1016/j.enbuild.2022.112164.

[27] H. Karimianfard, "A robust optimization framework for smart home energy management: Integrating photovoltaic storage, electric vehicle charging, and demand response," *Journal of Energy Storage*, vol. 110, p. 115 259, 2025. DOI: 10.1016/j.est.2024.115259.

[28] U. Zafar, S. Bayhan, and A. Sanfilippo, "Home energy management system concepts, configurations, and technologies for the smart grid," *IEEE Access*, vol. 8, pp. 119 271–119 286, 2020. DOI: 10.1109/ACCESS.2020.3005244.

[29] Y. O. Assolami, "Residential energy management systems: A comprehensive review and case study," *Engineering Science and Technology, an International Journal*, vol. 78, p. 102 355, 2026. DOI: 10.1016/j.jestch.2026.102355.

[30] Ö. Okur, P. Heijnen, and Z. Lukszo, "Aggregator's business models in residential and service sectors: A review of operational and financial aspects," *Renewable and Sustainable Energy Reviews*, vol. 139, p. 110 702, 2021. DOI: 10.1016/j.rser.2020.110702.

[31] X. Lu, K. Li, H. Xu, F. Wang, Z. Zhou, and Y. Zhang, "Fundamentals and business model for resource aggregator of demand response in electricity markets," *Energy*, vol. 204, p. 117 885, 2020. DOI: 10.1016/j.energy.2020.117885.

[32] T. Robbiano, M. Fresia, M. Caliano, and S. Bracco, "Optimal energy management strategies for aggregators in renewable energy communities," *International Journal of Electrical Power & Energy Systems*, vol. 172, p. 111 085, 2025. DOI: 10.1016/j.ijepes.2025.111085.

[33] K. Hoth, B. Wiegel, T. Schug, and K. Fischer, "The energy aggregator problem: A holistic mixed-integer linear programming approach," *Sustainable Energy, Grids and Networks*, vol. 43, p. 101 754, 2025. DOI: 10.1016/j.segan.2025.101754.

[34] S. Woltmann and J. Kittel, "Development and implementation of multi-agent systems for demand response aggregators in an industrial context," *Applied Energy*, vol. 314, p. 118 841, 2022. DOI: 10.1016/j.apenergy.2022.118841.

[35] P. Sharma, H. Dutt Mathur, P. Mishra, and R. C. Bansal, "A critical and comparative review of energy management strategies for microgrids," *Applied Energy*, vol. 327, p. 120 028, 2022. DOI: 10.1016/j.apenergy.2022.120028.

[36] B. Han, Y. Zahraoui, M. Mubin, S. Mekhilef, M. Seyedmahmoudian, and A. Stojcevski, "Home energy management systems: A review of the concept, architecture, and scheduling strategies," *IEEE Access*, vol. 11, pp. 19 999–20 025, 2023. DOI: 10.1109/ACCESS.2023.3248502.

[37] L. Strezoski, "Distributed energy resource management systems – DERMS: State of the art and how to move forward," *WIREs Energy and Environment*, vol. 12, no. 1, e460, 2023. DOI: 10.1002/wene.460.

[38] P. Pawar, Mudige TarunKumar, and Panduranga Vittal K., "An IoT based intelligent smart energy management system with accurate forecasting and load strategy for renewable generation," *Measurement*, vol. 152, p. 107 187, 2020. DOI: 10.1016/j.measurement.2019.107187.

[39] R. J. Bessa et al., "Data economy for prosumers in a smart grid ecosystem," in *Proceedings of the Ninth International Conference on Future Energy Systems*, ser. e-Energy '18, Association for Computing Machinery, 2018, pp. 622–630. DOI: 10.1145/3208903.3210282.

[40] G. S. Thirunavukkarasu, M. Seyedmahmoudian, E. Jamei, B. Horan, S. Mekhilef, and A. Stojcevski, "Role of optimization techniques in microgrid energy management systems: A review," *Energy Strategy Reviews*, vol. 43, p. 100 899, 2022. DOI: 10.1016/j.esr.2022.100899.

[41] D. Saha, N. Bazmohammadi, J. C. Vasquez, and J. M. Guerrero, "Multiple microgrids: A review of architectures and operation and control strategies," *Energies*, vol. 16, no. 2, p. 600, 2023. DOI: 10.3390/en16020600.

[42] S. R. Hukampal, P. Jena, and D. K. Singh, "Review of centralized, decentralized and hierarchical energy management system of microgrid," in *2025 1st IEEE Uttar Pradesh Section Women in Engineering International Conference on Electrical Electronics and Computer Engineering (UPWIECON)*, 2025, pp. 607–613, ISBN: 979-8-3315-6628-9. DOI: 10.1109/UPWIECON67212.2025.11390336.

[43] J. A. Rodriguez-Gil et al., "Energy management system in networked microgrids: An overview," *Energy Systems*, vol. 17, pp. 315–346, 2026. DOI: 10.1007/s12667-024-00676-6.

[44] H. M. Yassim, M. N. Abdullah, C. K. Gan, and A. Ahmed, "A review of hierarchical energy management system in networked microgrids for optimal inter-microgrid power exchange," *Electric Power Systems Research*, vol. 231, p. 110 329, 2024. DOI: 10.1016/j.epsr.2024.110329.

[45] R. Bustos, L. G. Marín, A. Navas-Fonseca, L. Reyes-Chamorro, and D. Sáez, "Hierarchical energy management system for multi-microgrid coordination with demand-side management," *Applied Energy*, vol. 342, p. 121 145, 2023. DOI: 10.1016/j.apenergy.2023.121145.

[46] International Renewable Energy Agency, "Demand-side flexibility for power sector transformation," International Renewable Energy Agency, Abu Dhabi, Analytical Brief, 2019. [Online]. Available: https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Dec/IRENA_Demand-side_flexibility_2019.pdf.

[47] J. Villar, R. Bessa, and M. Matos, "Flexibility products and markets: Literature review," *Electric Power Systems Research*, vol. 154, pp. 329–340, 2018. DOI: 10.1016/j.epsr.2017.09.005.

[48] G. Morales-España, R. Martínez-Gordón, and J. Sijm, "Classifying and modelling demand response in power systems," *Energy*, vol. 242, p. 122 544, 2022. DOI: 10.1016/j.energy.2021.122544.

[49] T. Freire-Barceló, F. Martín-Martínez, and Á. Sánchez-Miralles, "A literature review of explicit demand flexibility providing energy services," *Electric Power Systems Research*, vol. 209, p. 107 953, 2022. DOI: 10.1016/j.epsr.2022.107953.

[50] C. Cabot and M. Villavicencio, "The demand-side flexibility in liberalised power market: A review of current market design and objectives," *Renewable and Sustainable Energy Reviews*, vol. 201, p. 114 643, 2024. DOI: 10.1016/j.rser.2024.114643.

[51] H. Li, Z. Wang, T. Hong, and M. A. Piette, "Energy flexibility of residential buildings: A systematic review of characterization and quantification methods and applications," *Advances in Applied Energy*, vol. 3, p. 100 054, 2021. DOI: 10.1016/j.adapen.2021.100054.

[52] B. Williams, D. Bishop, P. Gallardo, and J. G. Chase, "Demand side management in industrial, commercial, and residential sectors: A review of constraints and considerations," *Energies*, vol. 16, no. 13, p. 5155, 2023. DOI: 10.3390/en16135155.

[53] M. Ranaboldo et al., "A comprehensive overview of industrial demand response status in europe," *Renewable and Sustainable Energy Reviews*, vol. 203, p. 114 797, 2024. DOI: 10.1016/j.rser.2024.114797.

[54] D. Ramin, S. Spinelli, and A. Brusaferri, "Demand-side management via optimal production scheduling in power-intensive industries: The case of metal casting process," *Applied Energy*, vol. 225, pp. 622–636, 2018. DOI: 10.1016/j.apenergy.2018.03.084.

[55] E. Zanetti, D. Blum, H. Fu, C. Weyandt, M. Pritoni, and M. A. Piette, "Commercial building HVAC demand flexibility with model predictive control: Field demonstration and literature insights," *Energy and Buildings*, vol. 345, p. 116 097, 2025. DOI: 10.1016/j.enbuild.2025.116097.

[56] P. A. D. Turner et al., "Embedding energy flexibility capability in air source heat pumps via third-party control: Insights from a field trial on residential buildings in england," *Applied Energy*, vol. 389, p. 125 705, 2025. DOI: 10.1016/j.apenergy.2025.125705.

[57] J. Bender, L. Fait, and H. Wetzel, "Acceptance of demand-side flexibility in the residential heating sector: Evidence from a stated choice experiment in germany," *Energy Policy*, vol. 191, p. 114 145, 2024. DOI: 10.1016/j.enpol.2024.114145.

[58] H. Johra, A. Marszal-Pomianowska, J. R. Ellingsgaard, and M. Liu, "Building energy flexibility: A sensitivity analysis and key performance indicator comparison," *Journal of Physics: Conference Series*, vol. 1343, no. 1, p. 012 064, 2019. DOI: 10.1088/1742-6596/1343/1/012064.

[59] N. Hekmat, H. Cai, T. Zufferey, G. Hug, and P. Heer, "Data-driven demand-side flexibility quantification: Prediction and approximation of flexibility envelopes," in *2023 IEEE Belgrade PowerTech*, 2023. DOI: 10.1109/POWERTECH55446.2023.10202703.

[60] F. D'Ettorre et al., "Exploiting demand-side flexibility: State-of-the-art, open issues and social perspective," *Renewable and Sustainable Energy Reviews*, vol. 165, p. 112 605, 2022. DOI: 10.1016/j.rser.2022.112605.

[61] H. Li et al., "Data-driven key performance indicators and datasets for building energy flexibility: A review and perspectives," *Applied Energy*, vol. 343, p. 121 217, 2023. DOI: 10.1016/j.apenergy.2023.121217.

[62] H. Li and T. Hong, "On data-driven energy flexibility quantification: A framework and case study," *Energy and Buildings*, vol. 296, p. 113 381, 2023. DOI: 10.1016/j.enbuild.2023.113381.

[63] X. Zhou et al., "Demand flexibility characterisation in non-residential buildings: A review," *Applied Energy*, vol. 393, p. 126 100, 2025. DOI: 10.1016/j.apenergy.2025.126100.

[64] A. Martín-Crespo, E. Baeyens, S. Saludes-Rodil, and F. Frechoso-Escudero, "Aggregated demand flexibility prediction of residential thermostatically controlled loads and participation in electricity balance markets," *International Journal of Energy Research*, vol. 2025, no. 1, p. 8 819 201, 2025. DOI: 10.1155/er/8819201.

[65] J. Park et al., "Probabilistic forecast-based portfolio optimization of electricity demand at low aggregation levels," *Applied Energy*, vol. 353, p. 122 109, 2024. DOI: 10.1016/j.apenergy.2023.122109.

[66] K. Li, Z. Li, C. Huang, and Q. Ai, "Online transfer learning-based residential demand response potential forecasting for load aggregator," *Applied Energy*, vol. 358, p. 122 631, 2024. DOI: 10.1016/j.apenergy.2024.122631.

[67] A. Petrucci, F. K. Ayevide, A. Buonomano, and A. Athienitis, "Development of energy aggregators for virtual communities: The energy efficiency-flexibility nexus for demand response," *Renewable Energy*, vol. 215, p. 118 975, 2023. DOI: 10.1016/j.renene.2023.118975.

[68] European Union Agency for the Cooperation of Energy Regulators. "Market rules for different electricity market timeframes." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.acer.europa.eu/electricity/market-rules/market-rules-different-electricity-market-timeframes.

[69] European Network of Transmission System Operators for Electricity. "Electricity balancing." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.entsoe.eu/major-projects/network-code-implementation/cross-border-electricity-balancing-pilot-projects/Pages/default.aspx.

[70] European Parliament and the Council of the European Union, *Regulation (EU) 2024/1747 of the european parliament and of the council of 13 june 2024 amending regulations (EU) 2019/942 and (EU) 2019/943 as regards improving the union's electricity market design*, 2024. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1747/oj/eng.

[71] OMIE, "Day-ahead and intraday electricity market rules," OMIE, Tech. Rep., 2025, Non-binding translation. CNMC Resolution DCOOR/DE/009/24. Accessed: Jul. 29, 2026. [Online]. Available: https://www.omie.es/sites/default/files/2025-03/r_28022025_en.pdf.

[72] NEMO Committee. "Single day-ahead coupling (SDAC)." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.nemo-committee.eu/sdac.

[73] European Commission, *Commission regulation (EU) 2015/1222 of 24 july 2015 establishing a guideline on capacity allocation and congestion management*, 2015. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2015/1222/oj/eng.

[74] Redes Energéticas Nacionais. "MIBEL: OMIE market results." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://mercado.ren.pt/EN/Electr/MarketInfo/MarketResults/OMIE.

[75] OMIE. "Day-ahead price." Official market data webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.omie.es/en/market-results/daily/daily-market/day-ahead-price.

[76] Entidade Reguladora dos Serviços Energéticos, "Analysis of portugal-spain interconnection capacity and monitoring of compliance with the minimum levels of available capacity for cross-zonal trade in 2024," Entidade Reguladora dos Serviços Energéticos, Tech. Rep., 2025. Accessed: Jul. 29, 2026. [Online]. Available: https://www.erse.pt/media/fvgbyebz/final-report-on-maczt-2024-interconnection-analysis-and-monitoring.pdf.

[77] European Commission. "EU electricity trading in the day-ahead markets becomes more dynamic." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://energy.ec.europa.eu/news/eu-electricity-trading-day-ahead-markets-becomes-more-dynamic-2025-10-01_en.

[78] European Parliament and Council of the European Union, *Regulation (EU) 2019/943 of the european parliament and of the council of 5 june 2019 on the internal market for electricity*, 2019. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2019/943/oj/eng.

[79] OMIE. "OMIE successfully implements trading in 15-minute intervals in the day-ahead market." Press release, Accessed: Jul. 29, 2026. [Online]. Available: https://www.omeldiversificacion.es/sites/default/files/2025-10/2501001_go_live_mtu15_md_en_vf.pdf.

[80] Entidade Reguladora dos Serviços Energéticos. "Understanding the invoice." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.erse.pt/en/energy-consumers/eletricity/understanding-the-invoice/.

[81] Entidade Reguladora dos Serviços Energéticos. "Tariffs and prices for electricity." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.erse.pt/en/activities/market-regulation/tariffs-and-prices-electricity/.

[82] Entidade Reguladora dos Serviços Energéticos. "What is included in the price." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.erse.pt/en/energy-consumers/eletricity/what-is-included-in-the-price/.

[83] European Commission, *Commission regulation (EU) 2017/1485 of 2 august 2017 establishing a guideline on electricity transmission system operation*, 2017. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2017/1485/oj/eng.

[84] European Commission, *Commission regulation (EU) 2017/2195 of 23 november 2017 establishing a guideline on electricity balancing*, 2017. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2017/2195/oj/eng.

[85] ENTSO-E. "PICASSO: Platform for the International Coordination of Automated Frequency Restoration and Stable System Operation." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.entsoe.eu/network_codes/eb/picasso/.

[86] ENTSO-E. "MARI: Manually Activated Reserves Initiative." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.entsoe.eu/network_codes/eb/mari/.

[87] ENTSO-E. “TERRE: Trans European Replacement Reserves Exchange.” Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.entsoe.eu/network_codes/eb/terre/.

[88] Sympower. “The differences in demand response - balancing services explained,” Accessed: Jul. 29, 2026. [Online]. Available: https://sympower.net/articles/the-differences-in-demand-response-balancing-services-explained.

[89] Replacement Reserve TSOs, *Announcement from replacement reserve TSOs*, 2024. Accessed: Jul. 29, 2026. [Online]. Available: https://eepublicdownloads.blob.core.windows.net/public-cdn-container/clean-documents/Network%20codes%20documents/NC%20EB/2024/Announcement_from_RR_TSOs.pdf.

[90] Redes Energéticas Nacionais. “End of TERRE.” Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://mercado.ren.pt/EN/Electr/Communication/News/Pages/News_20251211_1.aspx.

[91] ENTSO-E, *REN successfully accessed the MARI platform*, MARI TSOs communication note, 2024. Accessed: Jul. 29, 2026. [Online]. Available: https://eepublicdownloads.blob.core.windows.net/public-cdn-container/clean-documents/Network%20codes%20documents/NC%20EB/2024/241127_Communication_note_REN_go-live.pdf.

[92] Entidade Reguladora dos Serviços Energéticos, “Manual de procedimentos da gestão global do sistema do setor elétrico,” Entidade Reguladora dos Serviços Energéticos, Tech. Rep., 2025. Accessed: Jul. 29, 2026. [Online]. Available: https://www.erse.pt/media/q10chfti/mpggs_articulado-250911.pdf.

[93] ACER, *Decision no 11/2020 of 17 june 2020 on the methodology for a list of standard products for balancing capacity for frequency restoration reserves and replacement reserves*, 2020. Accessed: Jul. 29, 2026. [Online]. Available: https://www.acer.europa.eu/sites/default/files/documents/en/Electricity/MARKET-CODES/ELECTRICITY-BALANCING/17%20SPBC/Action%202%20-%20SPBC%20ACER%20decision.pdf.

[94] First Climate. “Green power purchase agreements (PPAs).” Corporate webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.firstclimate.com/leistungen/green-power-purchase-agreements-ppas?lang=en.

[95] C. Mittler, M. Bucksteeg, and P. Staudt, “Review and morphological analysis of renewable power purchasing agreement types,” *Renewable and Sustainable Energy Reviews*, vol. 211, p. 115 293, 2025. DOI: 10.1016/j.rser.2024.115293.

[96] K. Kapral, K. Soetaert, and R. Castro, “An off-site power purchase agreement (PPA) as a tool to protect against electricity price spikes: Developing a framework for risk assessment and mitigation,” *Energies*, vol. 17, no. 9, p. 2161, 2024. DOI: 10.3390/en17092161.

[97] R. Qorbanian, N. Löhndorf, and D. Wozabal, “Valuation of power purchase agreements for corporate renewable energy procurement,” *European Journal of Operational Research*, vol. 326, no. 3, pp. 530–543, 2025. DOI: 10.1016/j.ejor.2025.05.054.

[98] J. I. Peña, R. Rodríguez, and S. Mayoral, “Hedging renewable power purchase agreements,” *Energy Strategy Reviews*, vol. 55, p. 101 513, 2024. DOI: 10.1016/j.esr.2024.101513.

[99] ACER. "Power purchase agreements (PPAs)." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.acer.europa.eu/electricity/market-monitoring/ppas.

[100] European Commission, *Commission recommendation (EU) 2026/917 of 22 april 2026 on removing barriers to the development of power purchase agreements and other energy purchase agreements*, 2026. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/reco/2026/917/oj/eng.

[101] Presidência do Conselho de Ministros, *Decreto-lei n.º 99/2024, de 3 de dezembro*, 2024. Accessed: Jul. 29, 2026. [Online]. Available: https://diariodarepublica.pt/dr/detalhe/decreto-lei/99-2024-898705893.

[102] Ambiente e Energia, *Portaria n.º 367/2024/1, de 31 de dezembro*, 2024. Accessed: Jul. 29, 2026. [Online]. Available: https://diariodarepublica.pt/dr/detalhe/portaria/367-2024-901661930.

[103] Entidade Reguladora dos Serviços Energéticos, *Diretiva n.º 11/2025, de 18 de novembro*, Aprova o Manual de Procedimentos da Atividade de Registo e Contratação Bilateral de Energia Elétrica (MP PPA), 2025. Accessed: Jul. 29, 2026. [Online]. Available: https://diariodarepublica.pt/dr/detalhe/diretiva/11-2025-946776857.

[104] S. Kerscher and P. Arboleya, "The key role of aggregators in the energy transition under the latest european regulatory framework," *International Journal of Electrical Power & Energy Systems*, vol. 134, p. 107 361, 2022. DOI: 10.1016/j.ijepes.2021.107361.

[105] Y. Zahraoui, B. R. Pokhrel, I. Oleinikova, and S. N. Gros, "Assessment of TSO-DSO coordination to enhance flexible grid operation," *Current Sustainable/Renewable Energy Reports*, vol. 12, 2025, Art. no. 20. DOI: 10.1007/s40518-025-00269-6.

[106] European Parliament and Council of the European Union, *Directive (EU) 2019/944 of the european parliament and of the council of 5 june 2019 on common rules for the internal market for electricity and amending directive 2012/27/EU*, 2019. Accessed: Jul. 29, 2026. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2019/944/oj/eng.

[107] ACER, "Demand response and other distributed energy resources: What barriers are holding them back?" European Union Agency for the Cooperation of Energy Regulators, Tech. Rep., 2023. Accessed: Jul. 29, 2026. [Online]. Available: https://www.acer.europa.eu/barriers-demand-response-and-other-distributed-energy-resources-2023-market-monitoring-report.

[108] I. Saviuc, C. Lopez, A. Puskas, K. Rollert, and P. Bertoldi, "Explicit demand response for small end-users and independent aggregators: Status, context, enablers and barriers," Publications Office of the European Union, Luxembourg, Tech. Rep. EUR 31190 EN, 2022. DOI: 10.2760/625919. Accessed: Jul. 29, 2026. [Online]. Available: https://publications.jrc.ec.europa.eu/repository/handle/JRC129745.

[109] M. Otte, J. Kamsamrong, and S. Lehnhoff, "How can aggregators improve the TSO-DSO-customer coordination in digitalised power systems?" International Smart Grid Action Network, Tech. Rep., 2024. Accessed: Jul. 29, 2026. [Online]. Available: https://www.iea-isgan.org/wp-content/uploads/2024/07/ISGAN_Aggregator_final.pdf.

[110] ACER, "Unlocking flexibility: No-regret actions to remove barriers to demand response," European Union Agency for the Cooperation of Energy Regulators, Tech. Rep., 2025. Accessed: Jul. 29, 2026. [Online]. Available: https://www.acer.europa.eu/sites/default/files/documents/Publications/2025-ACER-Unlocking-flexibility-demand-response-barriers.pdf.

[111] K. Baltputnis, T. Schittekatte, and Z. Broka, "Independent aggregation in the nordic day-ahead market: What is the welfare impact of socializing supplier compensation payments?" *Heliyon*, vol. 11, no. 1, e41619, 2025. DOI: 10.1016/j.heliyon.2024.e41619.

[112] I. López et al., "European energy communities: Characteristics, trends, business models and legal framework," *Renewable and Sustainable Energy Reviews*, vol. 197, p. 114 403, 2024. DOI: 10.1016/j.rser.2024.114403.

[113] J. Iria and F. Soares, "An energy-as-a-service business model for aggregators of prosumers," *Applied Energy*, vol. 347, p. 121 487, 2023. DOI: 10.1016/j.apenergy.2023.121487.

[114] J. Iria, F. Soares, and M. Matos, "Optimal bidding strategy for an aggregator of prosumers in energy and secondary reserve markets," *Applied Energy*, vol. 238, pp. 1361–1372, 2019. DOI: 10.1016/j.apenergy.2019.01.191.

[115] C. Huang, K. Li, and N. Zhang, "Strategic joint bidding and pricing of load aggregators in day-ahead demand response market," *Applied Energy*, vol. 377, p. 124 552, 2025. DOI: 10.1016/j.apenergy.2024.124552.

[116] J. Cano-Martinez, A. Quijano-López, and V. Fuster-Roig, "A scoping review of flexibility markets in the power sector: Models, mechanisms, and business perspectives," *Energies*, vol. 18, no. 19, p. 5213, 2025. DOI: 10.3390/en18195213.

[117] L. Rodrigues, F. Coelho, J. Mello, and J. Villar, "Digital platforms to support the flexibility value chain, run flexibility markets, and manage energy communities," *Current Sustainable/Renewable Energy Reports*, vol. 12, 2025, Art. no. 15. DOI: 10.1007/s40518-025-00264-x.

[118] Viegand Maagøe, "Overview of flexibility platforms," EDNA - Efficient, Demand Flexible Networked Appliances, Report DF7, Jan. 21, 2025. Accessed: Jul. 29, 2026. [Online]. Available: https://www.iea-4e.org/wp-content/uploads/2025/01/Report-Flexibility-platforms-EDNA_01-FB.pdf.

[119] S. Potenciano Menci and O. Valarezo, "Decoding design characteristics of local flexibility markets for congestion management with a multi-layered taxonomy," *Applied Energy*, vol. 357, p. 122 203, 2024. DOI: 10.1016/j.apenergy.2023.122203.

[120] F. Dominguez et al., "A set of standardised products for system services in the TSO-DSO-consumer value chain," OneNet Project, Horizon 2020, Deliverable D2.2, 2021. DOI: 10.5281/zenodo.5925738. Accessed: Jul. 29, 2026. [Online]. Available: https://onenet-project.eu/wp-content/uploads/2022/10/D22-A-set-of-standardised-products-for-system-services-in-the-TSO-DSO-consumer-value-chain.pdf.

[121] EU DSO Entity, "Report on distributed flexibility practices: Markets for local services," EU DSO Entity, Tech. Rep., 2026. Accessed: Jul. 29, 2026. [Online]. Available: https://eudsoentity.eu/wp-content/uploads/2026/02/Report-on-Distributed-Flexibility-Practices-Markets-for-Local-Services-2.pdf.

[122] Piclo. "The marketplace for energy flexibility." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.piclo.com/.

[123] NODESmarket. "Transacting flexibility with cutting-edge technology." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://nodesmarket.com/.

[124] GOPACS. "Platform for congestion management." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.gopacs.eu/en/.

[125] E-REDES. "FIRMe: Integration of renewable energy." Official webpage, Accessed: Jul. 29, 2026. [Online]. Available: https://www.e-redes.pt/en/energy-transition/firme.

[126] F. Conte, M. Saviozzi, and S. Grillo, "An optimization problem for day-ahead planning of electrical energy aggregators," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 12 213–12 220, 2020. DOI: 10.1016/j.ifacol.2020.12.1098.

[127] J. P. Iria, F. J. Soares, and M. A. Matos, "Trading small prosumers flexibility in the energy and tertiary reserve markets," *IEEE Transactions on Smart Grid*, vol. 10, no. 3, pp. 2371–2382, 2019. DOI: 10.1109/TSG.2018.2797001.

[128] J. Iria, P. Scott, and A. Attarha, "Network-constrained bidding optimization strategy for aggregators of prosumers," *Energy*, vol. 207, p. 118 266, 2020. DOI: 10.1016/j.energy.2020.118266.

[129] S. Afzali, M. Parsa Moghaddam, M. K. Sheikh-El-Eslami, and R. Zamani, "A flexibility-oriented bidding strategy for virtual power plants incorporating local energy communities: A bi-level stochastic-robust methodology," *Applied Energy*, vol. 399, p. 126 355, 2025. DOI: 10.1016/j.apenergy.2025.126355.

[130] M. Vahid-Ghavidel, M. S. Javadi, S. F. Santos, M. Gough, M. Shafie-khah, and J. P. S. Catalão, "Energy storage system impact on the operation of a demand response aggregator," *Journal of Energy Storage*, vol. 64, p. 107 222, 2023. DOI: 10.1016/j.est.2023.107222.

[131] D. Fernández-Muñoz and J. I. Pérez-Díaz, "Optimisation models for the day-ahead energy and reserve self-scheduling of a hybrid wind-battery virtual power plant," *Journal of Energy Storage*, vol. 57, p. 106 296, 2023. DOI: 10.1016/j.est.2022.106296.

[132] G. Gousis, N. Koltsaklis, K. Oureilidis, and G. Christoforidis, "Enhancing the viability and bankability of hybrid RES-BESS systems with corporate power purchase agreements and electricity market participation," *Renewable Energy*, vol. 256, p. 124 631, 2026. DOI: 10.1016/j.renene.2025.124631.

[133] I. Rezende et al., "Seaport energy management system considering greenhouse gas emissions," *Smart Grids and Sustainable Energy*, vol. 11, no. 1, 2026, Art. no. 19. DOI: 10.1007/s40866-026-00337-w.

[134] I. Papias, E. Sarantinopoulos, V. Michalakopoulos, E. Sarmas, and V. Marinakis, "FLEXiT: A platform for estimating residential energy flexibility for aggregated and disaggregated demand-side management," *SoftwareX*, vol. 32, p. 102 433, 2025. DOI: 10.1016/j.softx.2025.102433.

[135] Z. Wang and H. Zhang, "Customer baseline load estimation for virtual power plants in demand response: An attention mechanism-based generative adversarial networks approach," *Applied Energy*, vol. 357, p. 122 544, 2024. DOI: 10.1016/j.apenergy.2023.122544.

[136] C. Tarmanini, N. Sarma, C. Gezegin, and O. Ozgonenel, "Short term load forecasting based on ARIMA and ANN approaches," *Energy Reports*, vol. 9, no. Supplement 3, pp. 550–557, 2023. DOI: 10.1016/j.egyr.2023.01.060.

[137] B. Magalhães, P. Bento, J. Pombo, M. R. Calado, and S. Mariano, "Short-term load forecasting based on optimized random forest and optimal feature selection," *Energies*, vol. 17, no. 8, p. 1926, 2024. DOI: 10.3390/en17081926.

[138] M. C. Ruiz-Abellón, L. A. Fernández-Jiménez, A. Guillamón, and A. Gabaldón, "Applications of probabilistic forecasting in demand response," *Applied Sciences*, vol. 14, no. 21, p. 9716, 2024. DOI: 10.3390/app14219716.

[139] J. Ribeiro, *Inputs and results dataset for aggregated EMS energy and reserve market planning of a B2C consumer portfolio - ATE project, portugal, 2022*, version 3, Mendeley Data, Jul. 12, 2026. DOI: 10.17632/3snn5fwmbs.3.

[140] Entidade Reguladora dos Serviços Energéticos. "Tarifas e preços regulados." Separador "Anos anteriores", ano de 2022, Accessed: Jul. 29, 2026. [Online]. Available: https://www.erse.pt/atividade/regulacao/tarifas-e-precos-eletricidade/?filtro=Anos%20anteriores#tarifas-e-precos-regulados.

[141] REN. "REN data hub," Accessed: Jul. 29, 2026. [Online]. Available: https://datahub.ren.pt/.

Appendix A

Detailed Consumer-Level Results

Table A.1: Share of energy purchased through each channel, per consumer, in Case C and Case F.

	Case C		Case F		
Consumer	**Tariff**	**Market**	**Tariff**	**Market**	**Bilateral**
User 1	76.2%	23.8%	72.0%	19.6%	8.4%
User 2	66.5%	33.5%	78.3%	21.3%	0.4%
User 3	71.4%	28.6%	67.5%	20.9%	11.6%
User 4	76.5%	23.5%	77.1%	22.7%	0.2%
User 5	57.2%	42.8%	49.0%	37.7%	13.3%

Table A.2: Annual variable cost per consumer across all evaluated configurations.

	Annual variable cost (€)				
Configuration	**User 1**	**User 2**	**User 3**	**User 4**	**User 5**
Baseline Simple	2202.7	268.0	3846.0	1632.9	1634.8
Baseline BESS	1956.5	257.1	3657.9	1550.2	1491.9
Case A	1089.1	257.2	2120.2	1549.7	1149.1
Case B	929.1	180.5	2104.9	1415.0	1017.4
Case C	834.9	134.7	1886.2	1297.8	834.7
Case D	−51.6	−756.4	1021.3	441.6	−32.6
Case E	−83.1	−773.5	984.6	386.3	−63.9
Case F	−121.2	−770.7	884.1	390.0	−97.7

Table A.3: Reserve contribution and attributed revenue per consumer in Case F.

Consumer	Contribution [kWh]	Band revenue [€]	Expected activation revenue [€]	Total [€]
User 1	13925.1	606.4	326.9	933.2
User 2	13558.1	596.7	316.0	912.7
User 3	13754.5	602.1	322.5	924.6
User 4	13755.0	600.8	303.0	903.8
User 5	13695.3	601.0	317.1	918.1

Table A.4: SSR per consumer across all evaluated configurations.

Configuration	User 1	User 2	User 3	User 4	User 5
Baseline Simple	14.0%	26.5%	3.6%	16.9%	15.9%
Baseline BESS	16.6%	30.1%	3.9%	22.1%	22.3%
Case A	24.0%	30.0%	4.0%	22.1%	26.8%
Case B	15.0%	26.2%	3.6%	16.6%	20.7%
Case C	13.6%	25.5%	3.4%	15.8%	18.2%
Case D	13.8%	23.6%	3.4%	15.4%	18.5%
Case E	13.9%	22.8%	3.3%	15.6%	16.8%
Case F	13.6%	23.0%	3.3%	15.9%	16.4%

Table A.5: SCR per consumer across all evaluated configurations.

Configuration	User 1	User 2	User 3	User 4	User 5
Baseline Simple	48.1%	88.0%	89.9%	64.7%	52.1%
Baseline BESS	56.9%	100.0%	98.6%	84.6%	73.1%
Case A	82.1%	99.9%	100.0%	84.6%	87.8%
Case B	51.1%	87.1%	91.1%	63.7%	67.7%
Case C	46.4%	84.7%	85.6%	60.7%	59.5%
Case D	47.1%	78.4%	85.4%	59.1%	60.5%
Case E	47.5%	75.8%	82.7%	59.6%	55.1%
Case F	46.6%	76.5%	82.4%	60.9%	53.6%