POLITECNICO DI TORINO

MASTER's Degree in ICT for Smart Societies

MASTER's Degree Thesis

Quantifying the Value of Residential Flexibility and Technology Upgrades: A Scenario-Ladder Framework with CAPEX and OPEX Optimization Analysis

Supervisors

Prof. Tao HUANG

Prof. Seyedmahmood HOSSEINIIMANI

Candidate

Maryam BIGONAH

MARCH 2026

Abstract

Residential buildings are central to the ongoing energy transition, in which improving energy efficiency, electrifying end uses, and integrating distributed renewable technologies have become key pathways for modernization. In this context, the performance of a residential energy system depends not only on the amount of energy consumed, but also on the timing of service demand and on the technologies used to provide it. However, practical and decision-oriented methods for comparing alternative upgrade pathways while preserving the same level of required household energy services remain limited.

To address this gap, this thesis develops a service-preserving techno-economic optimization framework for the evaluation of residential energy-system upgrades. The proposed methodology represents household demand in terms of required services and evaluates alternative modernization pathways through a scenario ladder. The optimization is organized to isolate (i) baseline cost accounting, (ii) shared rooftop photovoltaic generation, (iii) shared rooftop photovoltaic generation with shared battery dispatch, (iv) shared rooftop photovoltaic generation and shared battery dispatch combined with the scheduling of shiftable service-tasks, and (v) a separate thermal-retrofit comparison between an air-source heat pump and a gas boiler under an annualized CAPEX+OPEX evaluation. The framework is applied to a multi-unit residential case study in Turin, Italy, at hourly resolution.

The results show that the framework clearly separates the contribution of each intervention layer. For the analyzed case study, shared photovoltaic integration provides the largest initial reduction in electricity cost, shared battery storage yields additional savings through intertemporal energy shifting, and the scheduling of shiftable service-tasks provides a further incremental benefit. In the thermal-retrofit comparison, the gas boiler is preferred over the air-source heat pump under the adopted assumptions, since the lower operating cost of the heat pump is not sufficient to offset its higher annualized capital cost. Overall, the thesis provides a transparent decision-support framework for assessing residential refurbishment and electrification pathways under unchanged service requirements.

Keywords: Residential Energy Systems, Mixed-Integer Linear Programming, Techno-Economic Optimization, Shared Photovoltaic Generation, Shared Battery Energy Storage, Thermal Retrofit.

Acknowledgments

I would like to express my sincere gratitude to my supervisors, **Professor Tao Huang** and **Professor Seyedmahmood Hosseiniimani**, for their valuable guidance, constructive feedback, and continuous support throughout the development of this thesis. Their academic insight, encouragement, and patience were essential to the completion of this work and greatly enriched my learning experience.

My deepest gratitude goes to my parents, whose unconditional love, endless patience, and countless sacrifices are the true foundation of this achievement. In a time of separation and profound hardship, their unwavering belief in me has been my anchor, giving me the strength and resilience to overcome every challenge and reach this milestone.

Finally, I would like to thank everyone who supported me, directly or indirectly, during this journey. Their kindness, encouragement, and presence are sincerely appreciated.

Table of Contents

List of Figures

List of Tables

Acronyms

ARERA	Autorità di Regolazione per Energia Reti e Ambiente
ASHP	Air-Source Heat Pump
BESS	Battery Energy Storage System
CAPEX	Capital Expenditure
COP	Coefficient of Performance
CRF	Capital Recovery Factor
DHI	Diffuse Horizontal Irradiance
DHW	Domestic Hot Water
DNI	Direct Normal Irradiance
EV	Electric Vehicle
GHI	Global Horizontal Irradiance
GTI	Global Tilted Irradiance
HEMS	Home Energy Management System
ICT	Information and Communication Technology
LPG	LoadProfileGenerator
MAE	Mean Absolute Error
MILP	Mixed-Integer Linear Programming
O&M	Operation and Maintenance
OPEX	Operating Expenditure
Open-Meteo	Open-Meteo Weather Data Service
PV	Photovoltaic
PVGIS	Photovoltaic Geographical Information System
RMSE	Root Mean Squared Error
SOC	State of Charge
TOU	Time-of-Use
XGBoost	Extreme Gradient Boosting

Chapter 1

Introduction

1.1 Background and context

Residential buildings are becoming one of the most important arenas of the energy transition. Their role is no longer limited to passive energy consumption, because building refurbishment, end-use electrification, distributed renewable generation, energy storage, and digital control are progressively changing how energy is produced, exchanged, and used at household level. In this context, residential modernization is not simply a matter of reducing annual energy demand, but of redesigning the way everyday services are delivered under new technical, economic, and environmental constraints [1, 2, 3].

At a broad level, this thesis is situated within the wider European transition toward decarbonization, electrification, and higher operational flexibility in the residential sector. This transition has several connected dimensions. First, existing buildings are under increasing pressure to reduce fossil-fuel dependence and improve efficiency. Second, households are progressively exposed to more interactive technologies such as rooftop photovoltaic (PV) systems, battery storage, smart meters, heat pumps, and electric-vehicle (EV) charging. Third, the economics of residential energy use is becoming increasingly time-sensitive, because the value of electricity depends not only on how much is consumed, but also on when it is generated, stored, imported, exported, or used [1, 4, 5]. As a result, the residential sector is moving away from a static view of energy use toward a more dynamic and coordinated view in which technologies, user routines, and tariff structures interact.

This transition is especially relevant for studies that aim to evaluate modernization pathways in a realistic and decision-oriented manner. The issue is no longer limited to annual household bills or total energy balances alone, but extends to the performance of the complete residential energy system and its ability to deliver required services at lower cost and with lower dependence on the grid or fossil fuels [1, 6]. In this sense, residential modernization should not be interpreted as a purely technical upgrade of devices, but as a broader reconfiguration of how energy-related needs are satisfied within buildings under changing technical and economic conditions.

A central implication of this transition is that residential energy demand should

not be interpreted as demand for electricity or gas in itself. The energy-services literature has shown clearly that what households actually value are the outcomes delivered through energy use, such as thermal comfort, domestic hot water, cooked meals, clean clothes, lighting, appliance use, and mobility. Energy is therefore an input, while the final object of demand is the service obtained from it [7, 8]. This distinction is not only conceptual; it has direct methodological consequences. If the analyst evaluates a new technology, a new operating rule, or a new control strategy without preserving the same service outcomes, the comparison becomes misleading. A household that appears to "save energy" by reducing comfort, postponing hot-water availability, or failing to complete a required appliance cycle is not necessarily improving its energy system in any meaningful decision-oriented sense [7, 8].

For this reason, a service-preserving viewpoint is particularly appropriate for residential refurbishment and control studies. It makes explicit what must remain unchanged when alternatives are compared. In the context of residential modernization, this means that options involving PV, storage, flexible electrical tasks, and thermal technology choices should be judged under the condition that the same household services are still delivered. The relevance of this perspective is especially strong in buildings, where the largest shares of final energy demand are linked to service-based needs such as space heating and domestic hot water rather than discretionary electrical consumption alone [2, 3, 9].

The real-world problem follows directly from this transition. Households and building owners are increasingly confronted with a growing set of technology and operation choices, but they usually lack a structured basis for comparing them. In practical terms, residential users may need to decide whether to install PV, whether battery storage is economically justified, whether loads should be shifted in time, whether a conventional gas boiler should be retained or replaced by an air-source heat pump, and how these choices interact over the course of a year. However, the economic attractiveness of these options is highly case-dependent and strongly influenced by timing, load shape, user routines, technology coupling, and seasonal demand variation [10, 9, 1].

This problem is made harder by the fact that many existing studies either simplify household demand excessively or evaluate technologies in a way that is not fully consistent with real residential decision-making. Some studies focus only on electricity flows and do not make explicit which services are preserved. Others optimize short representative periods and therefore miss seasonal effects, especially for heating-related services. Others still treat flexibility as a generic shiftable load, without distinguishing between non-shiftable demand, non-interruptible appliance cycles, interruptible charging processes, and thermal-service constraints. As a result, an apparently good techno-economic outcome may partly reflect modelling assumptions that are too permissive from the household point of view [5, 4, 11].

The issue is especially important in multi-unit residential settings, where heterogeneity of occupants and aggregated demand can create both opportunities and distortions. The timing diversity across households can improve the value of shared

resources such as PV and batteries, but it can also hide which gains come from technology and which come from unrealistic aggregation assumptions. Recent work on residential systems has repeatedly shown that profitability and operational value vary substantially across households with different load profiles, even when annual demand is similar [10, 1]. Therefore, a decision-support framework that ignores heterogeneity or uses only a stylized average household can provide weak guidance for real buildings.

From a user-oriented perspective, the problem can be expressed simply: households may have access to more flexible and cleaner technologies than before, but they do not automatically know which combination is worth adopting, how these technologies should be coordinated, or whether the expected savings remain valid once comfort and service requirements are protected. This is precisely the type of problem that motivates a service-preserving techno-economic assessment of residential modernization [10, 12].

This problem is important now because several structural trends are converging at the same time. The first is the growing pressure to decarbonize residential buildings, especially through the electrification of thermal uses. Heating remains one of the most difficult end uses to decarbonize, yet it is also one of the most important in residential systems. Recent literature on building demand and residential heating emphasizes that thermal services dominate the physical and economic relevance of residential energy systems, and that any serious modernization pathway must therefore consider space heating and domestic hot water explicitly rather than treating them as secondary details [2, 3, 9].

The second is that timing has become economically decisive. In a conventional flat-tariff setting, many household decisions could be approximated by annual energy balances. That is no longer sufficient when PV generation is concentrated around daylight hours, many household demands are concentrated in morning and evening peaks, and flexible technologies such as batteries, EVs, and heat pumps can move energy across time. Under these conditions, the same annual energy use can produce different annual costs depending on its hourly distribution [13, 14, 12]. What matters is not only total demand, but also the coincidence between demand, on-site generation, and price signals.

The third reason is that consumer response remains imperfect. The literature does not support the simplistic assumption that households will always respond optimally to price signals or automatically adopt the best tariff and technology bundle. Empirical evidence from Northern Italy shows that time-of-use tariffs can produce measurable load shifting, but not necessarily in a uniformly beneficial way, especially for evening demand peaks [14]. Experimental evidence likewise shows that consumers may struggle to understand and choose time-varying tariffs appropriately, and that complexity itself can become a barrier to beneficial participation [12]. Consequently, the existence of theoretical flexibility does not guarantee that it will translate into robust household benefit without explicit modelling and decision support.

The fourth reason is that residential technologies are becoming more tightly coupled. The profitability of PV depends on self-consumption opportunities; the value of batteries depends on both tariff structure and load timing; the attractiveness of heat pumps depends on electricity prices, building demand, and possible interaction with on-site generation; and EV charging introduces a new, potentially large but schedulable electrical load. In other words, the marginal value of one technology often depends on whether another technology is already present. This interdependence is one of the main reasons why isolated one-technology comparisons are often insufficient [9, 6, 1].

Recent Italian research reinforces this point in the context of shared residential systems. Studies on PV-based Italian energy communities show that flexible demand can improve self-consumption and reduce energy costs, but also that these benefits depend on user willingness to modify habits and on how discomfort or feasibility is represented in the model [15]. This is particularly relevant for the present thesis because it confirms that flexibility has value, but only under a modelling structure that treats household behavior and service requirements with sufficient realism.

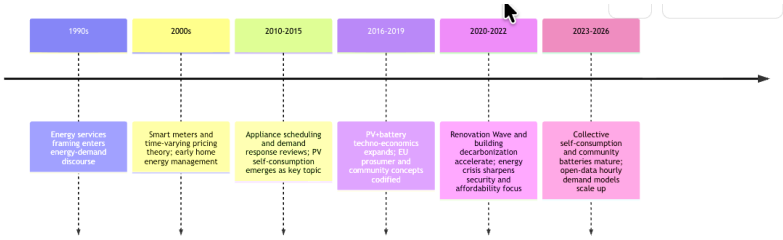

Figure 1.1: Research and policy evolution in residential energy modernization.

Figure 1.1 summarizes this evolution at a high level. The field has progressed from early household coordination concepts toward PV self-consumption, storage-enabled flexibility, prosumer participation, and, more recently, building-focused renovation and electrification agendas. This progression helps explain why residential modernization can no longer be treated as a static energy-efficiency problem, but instead requires a service-preserving and time-sensitive analytical framework.

Given these conditions, the problem should be framed as one of *service delivery under time-varying constraints*, not simply one of minimizing raw electricity consumption. This framing has several advantages. First, it keeps the analysis aligned with what households actually care about. A building is not modernized successfully merely because its imported electricity decreases; it is modernized successfully if it can still provide comfort, hot water, cleaning, cooking, and mobility services while lowering total cost, improving flexibility, or reducing fossil dependence [7, 8].

Second, a service-based framing naturally supports the distinction between different classes of demand. Some loads are effectively non-shiftable, because the service must be delivered at the time it occurs. Others are shiftable but non-interruptible, because once they start they must finish continuously. Others are interruptible but deadline-constrained, such as EV charging. Thermal services introduce yet another

layer, because part of their flexibility depends on technology type, thermal storage, or the inertia of the building itself [5, 9, 4]. This means that flexibility should not be treated as a single homogeneous resource. It should be represented as a structured set of service-delivering tasks and thermal demands with different operational rules.

Third, the service-based framing is essential for fair scenario comparison. One of the risks in residential optimization is obtaining cost savings by implicitly reducing welfare rather than improving system design. If comfort bands, appliance completion, hot-water availability, and charging requirements are not enforced, the model may report an optimal result that is numerically valid but decision-wise invalid. By contrast, when service levels are explicitly preserved, any economic improvement can be interpreted more credibly as the consequence of better timing, better technology choice, or better coordination among system components [7, 8, 11].

A further implication of this framing is the need for annual, high-resolution analysis. If residential modernization is evaluated under time-varying tariffs, PV production, seasonal heating demand, and heterogeneous occupant behavior, then a short representative-day approach is often insufficient. Work on domestic PV-battery systems has shown that time resolution can alter estimated self-consumption, financial outcomes, and even investment conclusions [11]. Likewise, large-scale analyses of residential load profiles have demonstrated that heterogeneity in household demand structure materially affects the profitability of PV-battery systems [10]. More broadly, recent building-demand modelling confirms that high-resolution temporal representations are necessary to capture the real interaction between weather, demand, and energy-system performance [2]. For the type of question addressed in this thesis, a full-year hourly representation is therefore not a luxury, but a methodological requirement.

The same reasoning also supports a scenario-based rather than a monolithic comparison. Since residential modernization involves several coupled decision layers, it is analytically useful to separate the incremental value of each layer. At a minimum, these layers include baseline cost accounting under existing service demand, on-site generation through shared PV, intertemporal shifting through shared battery dispatch, operational flexibility through the scheduling of shiftable tasks, and longer-term technology choice for thermal services. Such a stepwise structure improves interpretability because it shows where value is created, instead of collapsing all improvements into a single opaque co-optimization result [9, 3, 1].

summary

Against this background, the present thesis develops a service-preserving techno-economic framework for a residential case study in Italy. The analysis adopts a scenario-based structure in order to compare modernization pathways under unchanged household service requirements and to distinguish the contribution of generation, storage, flexibility, and thermal technology choices within one coherent decision framework. In this sense, the thesis responds directly to the current residential deci-

sion problem: not whether energy use can be reduced in the abstract, but whether a realistic residential system can be modernized in a way that is technically coherent, economically justified, and fair from the perspective of unchanged service delivery.

1.2 Literature review

Following the background and context outlined in the previous section, the relevant literature can be interpreted as an attempt to answer a central question: how can residential energy systems be modernized in a way that preserves required household services while improving techno-economic performance under time-varying conditions? In practice, researchers have approached this question through several connected strands. A first strand clarifies what should be modeled as the object of demand, namely services rather than electricity alone. A second strand develops Home Energy Management Systems (HEMS) and related optimization frameworks to coordinate distributed energy resources and household loads. A third strand studies photovoltaic (PV) self-consumption and battery dispatch under time-dependent tariffs. A fourth strand focuses on the scheduling of flexible tasks, especially non-interruptible appliance cycles and interruptible electric-vehicle charging. A fifth strand extends the analysis from short-term operation to longer-term technology choice through annualized techno-economic comparison. Taken together, these strands provide a strong methodological basis for residential decision support, while also revealing several limitations that directly motivate the present thesis [7, 8, 5, 4, 13, 1, 3].

A preliminary point concerns the data foundation on which residential optimization depends. Decision-oriented studies require temporally consistent inputs for demand, tariffs, weather, and PV generation, typically at hourly or sub-hourly resolution. When measured PV output is unavailable, researchers generally rely on measured inverter data, physics-based simulation from irradiance and temperature, or data-driven reconstruction based on meteorological variables [11, 16, 17]. This issue is relevant because the quality and temporal alignment of the PV series directly influence the estimated value of self-consumption, battery dispatch, and load shifting. The literature therefore makes clear that credible residential optimization begins not only with a suitable mathematical model, but also with coherent annual input data.

A first core stream of literature concerns the conceptual framing of residential demand. Fell argues that energy is better understood through the services it provides, and shows that the term “energy services” has often been used inconsistently across the literature [7]. Kalt et al. extend this idea through the Energy Service Cascade, which links energy carriers and technologies to services, benefits, and ultimately well-being [8]. For residential energy studies, this conceptual work is important because it shifts the modeling question from “how much electricity is consumed” to “how required household services are delivered.” This is especially relevant in buildings, where services such as thermal comfort, domestic hot water, cooking, cleaning, lighting, and mobility have different temporal characteristics and different degrees of flexibility. The main strength of this literature is that it provides a rigorous

conceptual basis for service-preserving comparison. Its main limitation, however, is that conceptual clarity does not by itself yield an operational optimization framework. In many applied studies, the language of services is adopted at a descriptive level, while the mathematical formulation still operates mainly on device-level electricity flows. As a result, the link between "flexible load" and "service-preserving flexibility" often remains implicit rather than explicitly enforced [7, 8].

This limitation leads directly to the HEMS literature, where researchers attempt to translate residential services into tractable decision variables and constraints. Reviews by Beaudin and Zareipour and by Paterakis et al. show that residential energy management has been addressed through a wide range of paradigms, including rule-based control, dynamic programming, metaheuristics, model predictive control, and mixed-integer optimization [5, 4]. Early HEMS studies such as Molderink et al. already emphasized the coordinated control of domestic generation, storage, and flexible demand under comfort-aware operation [18]. Over time, the field has increasingly converged toward formulations that treat the residential system as a combination of continuous energy-flow decisions and discrete service-related choices, particularly when appliances, storage, and local generation must be coordinated simultaneously.

Within this broader HEMS literature, mixed-integer linear programming (MILP) has remained one of the most widely adopted approaches. The reason is structural rather than accidental. Residential energy systems combine continuous variables such as grid import, battery state of charge, and PV allocation with discrete choices such as appliance start times, mutually exclusive charge/discharge states, and mode selection. MILP is well suited to this mixed structure while preserving transparency in the objective function and constraints [5, 4, 18]. Typical MILP formulations include an hourly power balance, battery dynamics with efficiency losses and state bounds, tariff-based operating costs, and binary decision variables for appliance cycles or operating modes [9, 19, 20]. One of the main strengths of this literature is therefore methodological maturity: it provides a clear and interpretable mathematical language for residential scheduling under realistic technical constraints.

The literature also explains why MILP remains particularly attractive in scenario-based studies such as the present thesis. In real-time control applications, model predictive control or heuristic methods may be preferable for computational reasons. However, when the research objective is transparent comparison across nested scenarios, MILP offers a major advantage: it makes the origin of savings easier to interpret. The analyst can observe how outcomes change when one decision layer is added, whether through PV, battery dispatch, flexible service scheduling, or technology replacement. In this sense, MILP is not only a solver choice, but also a way to preserve analytical clarity [5, 4, 1]. At the same time, the literature also reveals several limitations. MILP formulations can become computationally heavy over annual horizons, especially when many households, flexible devices, or detailed thermal representations are included. Moreover, comfort is sometimes represented through soft penalty terms that are difficult to calibrate consistently, which can

blur the distinction between genuine economic improvement and implicit service degradation [5, 4, 11].

A second major research strand concerns PV self-consumption and battery dispatch. Here the main problem is well established: PV generation is concentrated in daylight hours, whereas many residential demands are concentrated in the morning and evening. This temporal mismatch means that annual PV production alone is not a sufficient indicator of value; what matters is how much of that production is self-consumed, shifted, or exported, and under which tariff conditions [13, 10, 21]. Luthander et al. provide one of the most influential reviews in this area, showing that batteries and demand-side management can increase PV self-consumption, although the magnitude of the gain depends strongly on load profile, tariff structure, and export remuneration [13]. Schopfer et al. reinforce this point by demonstrating, across thousands of household load profiles, that profitability varies substantially even among households with similar annual demand [10]. Beck et al. reach a related conclusion when heat pumps, PV, and storage are considered jointly, showing that the value of one technology depends on the presence and operation of others [9].

This literature makes an important contribution by establishing both a physically meaningful operational framework and a set of useful performance indicators. Residential PV-battery studies commonly report self-consumption ratio, self-sufficiency, grid-import reduction, peak effects, and battery cycling metrics alongside cost outcomes [13, 10, 21]. These indicators are valuable because they distinguish energy-autonomy effects from purely tariff-driven economic effects. The literature also makes clear that conclusions are highly policy-dependent: when export remuneration is low relative to retail import prices, the model favors self-consumption and storage charging; when export compensation is more attractive, battery value can fall significantly [13, 21]. This is directly relevant for an Italy-focused residential case study, where the value of self-consumption depends on the adopted tariff and sharing assumptions.

From the perspective of the present thesis, however, this literature also has clear weaknesses. First, many PV-battery studies focus on dispatch quality or optimal sizing without explicitly addressing service preservation. Demand often enters as an exogenous load profile rather than as a set of service-delivering requirements. Second, the representation of the building is often limited to a single household or a representative load, which weakens the applicability of conclusions to multi-unit settings with occupant heterogeneity [10, 1]. Third, some studies co-optimize investment and operation in a single integrated model. While mathematically valid, this can make it difficult to separate the value of operational flexibility from the value of simply adding more equipment. For a decision-support thesis that aims to explain where benefits come from, this reduced interpretability is a significant limitation [1, 3].

A third literature stream concerns the scheduling of flexible demand. In this strand, researchers recognize that residential flexibility is heterogeneous rather than uniform. Some loads are essentially non-shiftable because the service must be delivered at the time it occurs. Others are shiftable but non-interruptible, because

once they begin they must continue until completion. Others, such as EV charging, are naturally interruptible but must satisfy an energy requirement by a deadline [5, 4]. This distinction has become standard in the scheduling literature because it maps more closely onto actual household services. Sou et al. model smart-appliance scheduling through MILP while explicitly representing uninterruptible and sequential operating constraints [20]. Bradač et al. present a MILP-based formulation for domestic appliance scheduling under realistic tariff structures and peak restrictions [19]. More recent work has extended this line by incorporating user dissatisfaction or customized preference modeling, showing that cost minimization alone is often insufficient to explain schedule acceptability [22].

The main strength of this literature is that it moves the field closer to actual service-based flexibility. Rather than assuming that demand can be shifted arbitrarily, these studies model tasks and charging sessions with durations, time windows, power limits, and completion rules [20, 19]. This is precisely the direction required for service-preserving residential optimization. It is also where HEMS studies become more credible from a user perspective, because they begin to distinguish between technically feasible schedules and behaviorally plausible ones. In addition, this literature shows that flexibility affects not only cost, but also local peaks and network stress. If many tasks are shifted into the same low-price period, unrealistic demand concentration may occur, which is why many formulations include explicit power caps or discomfort constraints [19, 4].

Important limitations remain, however. Many appliance-scheduling studies are short-horizon, often day-ahead or week-ahead, and therefore do not fully capture the interaction of flexibility with full-year seasonality. Thermal services are also frequently simplified or excluded. In many papers, the main flexible devices are washing machines, dishwashers, dryers, and EVs, while space heating and domestic hot water are either represented crudely or ignored [20, 19, 22]. This creates a mismatch with residential decision problems in which thermal services dominate annual energy use. Moreover, user comfort is often monetized through generic penalty terms rather than enforced as hard service-preserving constraints. Such formulations are useful for trade-off analysis, but are less appropriate when the aim is to compare modernization pathways under unchanged required services.

A fourth stream extends the literature from operation to techno-economic comparison and technology choice. Here, the question is no longer only how to dispatch energy or shift tasks, but how to compare alternative technologies that provide the same service. This is particularly important for thermal retrofit decisions such as boiler replacement, heat-pump adoption, and the combined assessment of PV, storage, and electrified heating [9, 1, 3, 23, 24]. Terlouw et al. provide a strong recent example of integrated residential energy-system design, showing how prospective costs and environmental indicators can be considered jointly for residential technology portfolios [1]. Demel et al. show the value of decision-support frameworks for heating decarbonization, emphasizing that residential technologies should be compared on a consistent service basis rather than on raw energy input alone [3]. Studies such

as Nicoletti et al. and Zhang et al. likewise show that once PV, storage, and heat pumps are considered together, the problem becomes one of joint sizing and annual operating cost rather than one of simple device substitution [23, 24].

This literature contributes two particularly important ideas. First, it provides formal tools to combine annualized CAPEX with OPEX, often through capital recovery factors or net-present-value logic, so that investment-stage and operating-stage costs can be expressed on a common annual basis [1, 3, 23]. Second, it reinforces the idea that residential technologies should be compared by the annual equivalent cost of delivering the same service. This is especially important in heat-pump versus boiler comparisons, where lower energy input for the heat pump does not automatically imply lower annual cost once capital expenditure and maintenance are included.

At the same time, the techno-economic literature also leaves several issues open. A recurring problem is the tendency to merge too many decision layers into a single co-optimization. Although mathematically elegant, this can make the results difficult to interpret for an end user. Another issue is that many studies optimize technology sizing with stylized demand or simplified behavioral assumptions, which weakens the credibility of decision-support conclusions for specific buildings [1, 23]. In thermal retrofit studies, COP is often simplified, building heat demand is treated as exogenous, or user comfort is assumed rather than verified. These assumptions are understandable for tractability, but they materially influence the economic outcome and should therefore be made explicit [9, 3]. Finally, much of the recent literature on forecast-aware optimization is moving toward renewable-energy-community settings with shared coordination mechanisms and forecast uncertainty, as shown by recent work such as Frieß et al. [17]. While valuable, that line of research does not fully resolve the behind-the-meter decision problem of a single multi-unit building treated as one entity, which is the focus of the present thesis.

Table 1.1 synthesizes the main studies reviewed above and highlights their contribution and remaining limitation from the perspective of this thesis.

Taken together, the reviewed literature already provides a substantial methodological basis for residential decision support. It offers a clear conceptual framing of energy services, mature optimization approaches for HEMS, physically grounded formulations for PV and battery dispatch, established models for scheduling shiftable appliance and EV demand, and formal methods for annualized techno-economic comparison. In that sense, the field is not lacking methods; rather, it lacks a sufficiently coherent combination of methods for the specific decision problem addressed in this thesis.

The main unresolved issue is not whether residential flexibility can be modeled, but whether it can be modeled in a way that remains transparently service-preserving and interpretable across multiple decision layers. In many studies, flexibility is still represented as generic shiftable demand rather than as clearly defined service-delivering tasks. In others, technology adoption and hourly scheduling are co-optimized so tightly that it becomes difficult to separate the value of PV, storage,

Study	Main focus	Contribution and remaining limitation
Beaudin and Zareipour (2015)	Review of HEMS modelling approaches	Summarizes the main paradigms for residential energy management and highlights modeling complexity. It does not, however, resolve how service preservation should be enforced in annual building-scale comparison.
Paterakis et al. (2017)	Demand response and flexibility classification	Reviews key demand-response formulations and international experience. Its limitation is that it is less focused on service-equivalent residential retrofit decisions.
Luthander et al. (2015)	PV self-consumption and demand-side management	Shows how batteries and DSM can increase PV self-consumption. Results, however, remain highly sensitive to tariff structure, export remuneration, and demand timing.
Schopfer et al. (2018)	Techno-economic assessment of PV-battery systems	Demonstrates strong heterogeneity in profitability across households. Service representation, however, remains implicit rather than explicitly task-based.
Sun et al. (2020)	Time-resolution effects in PV-battery modelling	Shows that timestep choice can materially change financial and self-consumption results, supporting the need for annual high-resolution analysis. It does not, however, address service-level modelling directly.
Sou et al. (2011) and Bradač et al. (2015)	MILP-based appliance scheduling	Provide explicit formulations for non-interruptible appliance cycles and tariff-based scheduling. Most applications, however, remain short-horizon and only partially integrate thermal services.
Beck et al. (2017)	Joint operation of PV, storage, and heat pumps	Demonstrates that technology value depends strongly on interaction among components. Results remain case-specific and do not isolate service-preserving scheduling as a separate decision layer.
Terlouw et al. (2023)	Integrated residential energy-system design	Provides a strong techno-economic framework with annualized investment logic. Joint optimization, however, can reduce interpretability regarding which intervention layer creates value.
Demel et al. (2023)	Heating decarbonization decision support	Strengthens the annual CAPEX+OPEX comparison logic for thermal technologies. It is less focused on combining this comparison with operational flexibility in one staged residential framework.
Recent forecast-aware studies (e.g., Nastić et al., 2024; Frieß et al., 2025)	Open-data PV forecasting and forecast-based coordination	Show that forecasting and coordination can materially affect renewable utilization and network outcomes. Most recent applications, however, move toward community-scale settings rather than the behind-the-meter multi-unit building problem addressed here.

Table 1.1: Selected literature on service-preserving residential energy modernization.

service scheduling, and thermal technology choice. In still others, annual temporal consistency, occupant heterogeneity, or thermal-service realism remain limited [10, 11, 1, 3, 17]. These limitations do not weaken the value of the literature; rather, they identify the gap to which the present thesis responds.

Summary

Accordingly, the present work builds on the literature in three specific ways. First, it adopts the service-based perspective not only as a conceptual framing, but also as a modeling rule: required services must remain unchanged across scenarios. Second, it uses a scenario-ladder structure to separate operational value from investment-stage conclusions, thereby improving interpretability. Third, it combines PV and battery dispatch, service-task scheduling, and a focused thermal CAPEX+OPEX comparison within one consistent annual framework for a residential case in Italy. The next section formalizes this unresolved need as the research gap and defines the objectives of the thesis.

1.3 Thesis contribution

Despite the substantial progress reviewed in the previous section, an important gap remains between existing residential energy-system studies and the type of decision support required by households or building owners. Many studies optimize device-level electricity profiles, PV self-consumption, or battery dispatch, but the delivered services are not always stated explicitly and the link between a so-called shiftable load and a service-delivering task often remains implicit. As a result, it is not always clear whether the reported savings preserve the same level of service, such as the same completion of laundry and dishwasher cycles, the same electric-vehicle charging requirement, and the same domestic hot-water and space-heating demand. In addition, many techno-economic studies co-optimize investment decisions and hourly operation in a single formulation. Although mathematically valid, this can obscure how much of the final benefit comes from operational scheduling itself and how much comes from simply adding or oversizing equipment. Recent forecast-aware formulations have also received increasing attention, but many of them are developed for energy-community settings with additional coordination mechanisms, whereas the behind-the-meter decision problem of a single multi-unit building treated as one entity is less directly addressed. Finally, open-data PV estimation methods are often evaluated mainly in terms of prediction accuracy, but are less often coupled end-to-end with annual techno-economic optimization in a way that shows how PV estimation uncertainty propagates into dispatch decisions, annual costs, and key performance indicators [10, 11, 1, 3, 17, 16].

These unresolved issues motivate the present thesis. From a practical perspective, residential users increasingly face decisions about technology upgrades and operational changes, but they do not necessarily have a transparent framework for judging whether

such changes are economically justified once the same required services are preserved. The problem is therefore not only how to reduce energy use or cost in the abstract, but how to compare modernization pathways in a way that remains service-preserving, interpretable, and relevant for an end-user decision question. In this context, the main research question of the thesis can be stated as follows:

> How can a multi-unit residential building in Italy be modernized through PV integration, battery storage, shiftable service-task scheduling, and thermal technology choice while preserving required household services and providing transparent annual techno-economic decision support?

To answer this question, the thesis develops and applies a service-based techno-economic optimization framework for a residential building in Italy treated as a single entity behind one electricity meter. The framework minimizes total annual cost while ensuring that the same service demand is satisfied, and it evaluates progressively more advanced configurations using 8,760-hour inputs for the year 2025. The central contribution is therefore not a single new technology model in isolation, but a structured decision-support framework that combines service-based demand representation, annual techno-economic optimization, and staged scenario comparison within one coherent methodology.

Research Objectives

The specific objectives of the thesis are as follows:

1. **Service-based demand modelling.** Represent residential demand in terms of required services and translate these services into energy requirements and operational tasks, distinguishing non-shiftable consumption from shiftable service-delivering tasks with realistic constraints such as time windows, non-interruptibility, and completion requirements. The modeled services include, at minimum, laundry, dishwashing, domestic hot water, space heating, cooking, and mobility charging.

2. **Construction of a consistent annual input dataset.** Build a coherent set of 8,760-hour inputs for the year 2025, including Italian time-of-use electricity prices with hourly F1/F2/F3 assignment, a gas price representation suitable for annual cost accounting, and an hourly PV generation series obtained through a weather-driven estimation approach aligned with PVGIS-consistent assumptions.

3. **Formulation of a service-preserving MILP scenario ladder.** Develop a sequence of mixed-integer linear programming models that progressively add decision layers: baseline cost accounting, PV and battery dispatch, scheduling of shiftable service-delivering tasks, and a final thermal-technology comparison under annualized CAPEX+OPEX logic.

4. **Quantification of annual value and decision support.** Quantify annual savings and key performance indicators across scenarios, including annual cost, grid imports, PV self-consumption, peak power, battery utilization, and task-completion indicators, and provide a decision-oriented recommendation regarding PV, battery, and thermal technology choice.

The scenario-ladder structure is a deliberate response to the research gap rather than a simple presentation choice. Its purpose is to preserve interpretability by separating the contribution of each intervention layer. The baseline scenario defines the reference condition against which all later changes are measured. The PV and battery stage isolates the operational value of local renewable generation and storage. The next stage adds the scheduling of shiftable service-delivering tasks, allowing the analysis to identify the incremental value of load shifting beyond PV and storage alone. Finally, the thermal comparison moves from operational scheduling to annualized technology choice. In this way, the framework avoids the common ambiguity of single-step co-optimization, where operational value and investment value can become mixed in one opaque result. The scenario ladder is therefore appropriate because it answers the end-user decision question progressively and transparently: what is the value of each additional intervention, under unchanged service requirements, and which modernization pathway is actually justified for the studied residential case [9, 1, 3].

Structure of the Dissertation

The dissertation is structured into seven chapters and three appendices. Chapter 1 presents the research background and context, reviews the relevant literature, and defines the main contribution of the thesis. Chapter 2 introduces the modeling framework adopted for the residential energy system, including local generation, storage systems, electricity demand, heating systems, and electric-vehicle charging demand.

Chapter 3 describes the optimization model, including the mathematical definition of sets, indices, parameters, decision variables, the objective function, and the constraint set. Chapter 4 presents the scenario-ladder framework and details the progressive evaluation of the baseline case, shared PV and battery operation, service scheduling, and thermal retrofit alternatives.

Chapter 5 focuses on data preparation and the study case, including service-demand profiles, thermal-demand data, mobility-related demand, electricity tariffs, natural gas prices, and rooftop PV sizing and generation. Chapter 6 reports the simulation results and discussion, with emphasis on PV forecasting, scenario setup, and the comparative techno-economic evaluation of the considered scenarios. Chapter 7 concludes the dissertation by summarizing the main findings, highlighting the key contributions of the work, and outlining directions for future research.

The dissertation is completed by three appendices. Appendix A provides the photovoltaic sizing inputs and the PV time-series generation workflow. Appendix B

reports supplementary data-preparation details, while Appendix C summarizes the component lifetime and replacement assumptions adopted in the economic assessment.

Summary

the research gap addressed by this thesis is not the absence of residential optimization methods in general, but the absence of a sufficiently transparent and service-preserving framework that can connect annual data preparation, operational scheduling, staged technology assessment, and decision-oriented interpretation for a single multi-unit residential building. The thesis responds to this gap by integrating those elements into one structured methodology and by using the scenario ladder to reveal, step by step, where value is created and how that value should be interpreted.

Chapter 2

Modeling of Home Appliances

2.1 Introduction

This chapter presents the mathematical modeling of the main technologies and end-use components considered in the residential energy system analyzed in this thesis. The purpose of this chapter is to describe, in a component-wise manner, the physical and operational behavior of the technologies that supply, store, and consume energy within the household. These component-level models constitute the technical basis of the integrated optimization framework developed in Chapter 3.

The residential energy system considered in this work includes photovoltaic (PV) generation as the local renewable electricity source, a battery energy storage system (BESS) as the electrical storage technology, residential electricity demand associated with household appliances and other domestic loads, and heating technologies used to satisfy the building thermal demand. Each component is introduced separately in order to define its main operating principles, governing variables, and mathematical relations.

The chapter is organized as follows. First, the PV generation model is introduced to describe the local renewable electricity production available to the household. Second, the battery storage model is presented, including the state-of-charge dynamics and the associated operating constraints. Third, the residential electricity demand is formulated, with particular attention to the distinction between non-shiftable and shiftable appliance loads. Finally, the thermal demand and the main heating technologies considered in this work, namely the gas boiler and the air-source heat pump, are described. Together, these models define the component representation of the residential energy system and provide the foundation for the mathematical optimization model presented in the next chapter.

2.2 Generation

Photovoltaic (PV) technology is widely adopted in residential energy systems as a decentralized and renewable electricity generation source. PV systems convert incident solar radiation directly into electrical energy through the photovoltaic effect in semiconductor materials. Due to declining technology costs, increasing policy support, and the growing role of distributed renewable generation, rooftop PV systems have become a key component of residential prosumer systems and local energy communities [25, 26].

In residential applications, PV systems are typically installed on building rooftops and operate as local electricity generation units that partially offset household electricity demand. Depending on the system configuration and the adopted operational strategy, the electricity produced by the PV array can be consumed directly by the household, stored in a battery energy storage system, or exchanged with the electricity grid. The integration of PV generation in residential buildings can therefore reduce grid dependence, increase renewable self-consumption, and lower the carbon intensity of household energy use [27, 28].

From a modeling perspective, the electrical power generated by a photovoltaic system depends primarily on the incident solar irradiance, the effective conversion efficiency of the PV modules, and the installed panel area. A simplified physical representation of the instantaneous PV power output at time step t can be written as:

$$P_{PV,t} = \eta_{PV}\, A_{PV}\, G_t \tag{2.1}$$

where $P_{PV,t}$ is the electrical power generated by the PV system at time step t (kW), η_{PV} is the effective conversion efficiency of the PV array, A_{PV} is the active panel area (m^2), and G_t is the solar irradiance incident on the panel surface at time step t ($\mathrm{kW/m}^2$).

In system-level optimization studies, PV generation is commonly represented through an available generation profile derived from meteorological data or solar simulation tools. Under this representation, the PV power output can be expressed as:

$$P_{PV,t} = P_{PV}^{rated}\, g_t \tag{2.2}$$

where P_{PV}^{rated} denotes the installed nominal PV capacity (kW) and g_t is the normalized generation factor at time step t, typically bounded between 0 and 1. The normalized factor captures the time-varying availability of solar energy due to irradiance and weather conditions.

To ensure physical consistency, the generated PV power is constrained by the available renewable resource:

$$0 \leq P_{PV,t} \leq P_{PV,t}^{avail} \tag{2.3}$$

where $P_{PV,t}^{avail}$ represents the available PV generation at time step t. This formulation is particularly useful in operational models where only the realizable renewable generation can be allocated to local demand, battery charging, or grid exchange.

The time index t refers to the discrete simulation interval. In this thesis, the modeling framework is based on hourly time steps in order to capture the temporal variability of solar generation and household electricity demand. The PV generation profile is treated as an exogenous input to the optimization model and interacts directly with both the battery storage system and the residential electrical load.

The inclusion of PV generation in residential energy systems provides several operational and economic benefits. First, locally produced solar electricity can directly supply household demand and reduce imported electricity from the grid. Second, when coordinated with storage and appliance scheduling, PV generation can increase renewable self-consumption and reduce electricity procurement costs. Finally, distributed PV contributes to the broader decarbonization of the residential energy sector by displacing fossil-based electricity generation [29, 26].

In the context of this thesis, the PV system is modeled as the primary local electricity generation source of the residential building. Its generation profile constitutes one of the main time-dependent inputs of the subsequent optimization model.

2.3 Storage Systems

Battery Energy Storage System (BESS)

Battery Energy Storage Systems (BESS) play a central role in modern residential energy systems, especially in buildings equipped with rooftop photovoltaic generation. A battery enables surplus electrical energy produced during periods of high renewable generation to be stored and later discharged when local generation is insufficient to meet demand. In this way, battery storage increases the self-consumption of locally generated PV electricity, reduces grid imports, and improves the coordination between electricity generation and demand [27, 28].

In residential applications, battery systems are typically coupled with rooftop PV installations to store surplus electricity during daytime hours and release it during periods of higher demand, particularly in the evening. This time-shifting capability enhances the economic value of local renewable generation and can reduce peak demand from the grid [26].

The operation of a battery is commonly described through its state of charge (SOC), which represents the amount of energy stored in the battery at a given time. The SOC evolves dynamically according to the charging and discharging processes. The battery energy balance can be expressed as:

$$SOC_{t+1} = SOC_t + \eta_c P_{charge,t} \Delta t - \frac{P_{discharge,t} \Delta t}{\eta_d} \tag{2.4}$$

where SOC_t is the state of charge of the battery at time step t (kWh), $P_{charge,t}$ is the battery charging power (kW), $P_{discharge,t}$ is the battery discharging power (kW),

η_c and η_d are the charging and discharging efficiencies, respectively, and Δt is the duration of the time step (h).

The battery energy content is limited by the minimum and maximum allowable storage levels:

$$SOC_{min} \leq SOC_t \leq SOC_{max} \tag{2.5}$$

where SOC_{min} and SOC_{max} denote the lower and upper bounds of the usable storage capacity. These limits reflect technical constraints and are also relevant for limiting accelerated battery degradation.

The charging and discharging powers are limited by the technical ratings of the battery inverter and the storage system itself:

$$0 \leq P_{charge,t} \leq P_{charge}^{max} \tag{2.6}$$

$$0 \leq P_{discharge,t} \leq P_{discharge}^{max} \tag{2.7}$$

where P_{charge}^{max} and $P_{discharge}^{max}$ denote the maximum charging and discharging powers of the BESS, respectively.

The initial state of the battery must also be defined:

$$SOC_0 = SOC^{init} \tag{2.8}$$

where SOC^{init} represents the initial energy stored in the battery at the beginning of the simulation horizon. For cyclic daily or annual analyses, it is also common to impose a terminal condition:

$$SOC_T = SOC_0 \tag{2.9}$$

where T denotes the final time step of the optimization horizon. This condition prevents the optimization from obtaining artificial benefits by ending the horizon with an unrealistically depleted or overcharged battery. If such a cyclic condition is not enforced in the final optimization model, Equation (2.9) can be omitted.

The battery may also be characterized through its usable energy capacity:

$$E_{BESS}^{use} = SOC_{max} - SOC_{min} \tag{2.10}$$

where E_{BESS}^{use} is the usable battery capacity (kWh). This quantity defines the effective storage range available for system operation.

Within the residential energy system, the BESS acts as an intermediate component between generation and demand. When photovoltaic production exceeds household electricity consumption, the surplus electricity can be diverted to battery charging. Conversely, during periods of insufficient PV generation or high electricity prices, the battery can discharge and supply part of the electrical load. In this thesis, the BESS is therefore modeled as an actively controlled component that supports the temporal

redistribution of locally generated electricity.

2.4 Residential Electricity Demand

Residential electricity demand represents the total electrical power consumed by household appliances and domestic devices within the residential building. In distributed energy systems, load modeling is essential because the temporal profile of household consumption determines the interaction between local generation, energy storage, and grid electricity imports. Accurate representation of residential demand is therefore crucial for assessing the performance of residential energy management strategies [30, 31].

Household electricity consumption typically includes lighting, refrigeration, kitchen appliances, information and communication devices, cleaning equipment, and miscellaneous plug loads. The operation of these appliances depends strongly on user behavior and daily activity patterns, which produce marked variations in load over time. Residential load profiles commonly exhibit morning and evening peaks associated with typical occupant routines [31].

At each time step t, the total residential electricity demand can be expressed as the aggregation of the individual appliance loads:

$$P_{load,t} = \sum_{i=1}^{N} P_{i,t} \tag{2.11}$$

where $P_{load,t}$ is the total residential electricity demand (kW), $P_{i,t}$ is the electrical power consumed by appliance i at time step t, and N is the total number of appliances considered in the household.

For modeling purposes, residential electricity demand can be decomposed into non-shiftable and shiftable components:

$$P_{load,t} = P_{ns,t} + P_{sh,t} \tag{2.12}$$

where $P_{ns,t}$ denotes the non-shiftable load and $P_{sh,t}$ denotes the shiftable load at time step t.

Non-shiftable loads correspond to appliances and devices that must operate immediately when service is required and therefore cannot be easily rescheduled. Typical examples include lighting, refrigeration, and consumer electronics that depend directly on occupant presence and comfort requirements. By contrast, shiftable loads are associated with appliances whose operating time can be moved within an admissible time window without compromising the service provided. Examples include washing machines, dishwashers, and similar domestic appliances [32, 33].

The total shiftable demand can be represented as the aggregation of the shiftable appliances:

$$P_{sh,t} = \sum_{j=1}^{M} P_{sh,j,t} \tag{2.13}$$

where $P_{sh,j,t}$ is the power consumption of shiftable appliance j at time step t, and M is the number of shiftable appliances.

For each shiftable appliance, the total energy required to complete a service cycle may be expressed as:

$$E_j = \sum_{t \in \mathcal{T}_j} P_{sh,j,t} \Delta t \tag{2.14}$$

where E_j is the electrical energy required by appliance j over its operating horizon and $\mathcal{T}_j$ denotes the admissible set of time steps within which the appliance can operate. This relation is useful for representing appliance service requirements while allowing time shifting of operation within a predefined time window.

The presence of shiftable loads creates the possibility of demand-side management through appliance scheduling. By moving the operation of certain loads toward periods of high PV generation or lower electricity prices, the residential system can improve renewable self-consumption and reduce total electricity procurement costs. In this thesis, however, Chapter 2 only introduces the component model of residential electricity demand; the actual scheduling decisions are formulated later within the optimization model.

2.5 Heating Systems

Thermal energy demand represents a major share of final energy consumption in residential buildings, particularly in regions where space heating is required during cold seasons. In addition to space heating, domestic hot water (DHW) demand contributes significantly to total household thermal consumption. The accurate modeling of thermal demand and heating technologies is therefore necessary for evaluating alternative heating solutions and for representing the interaction between thermal and electrical energy use in residential systems [34, 27].

In this thesis, two heating technologies are considered for the provision of useful heat to the household: a gas boiler and an air-source heat pump (ASHP). These systems represent two widely adopted alternatives in residential buildings and reflect the comparison between fossil-fuel-based and electrically driven heating solutions.

Thermal Demand Model

The total thermal demand of the household is assumed to be the sum of space heating demand and domestic hot water demand:

$$Q_{total,t} = Q_{SH,t} + Q_{DHW,t} \tag{2.15}$$

where $Q_{total,t}$ is the total thermal demand at time step t (kW), $Q_{SH,t}$ is the space heating demand, and $Q_{DHW,t}$ is the domestic hot water demand.

The thermal energy supplied by the selected heating technology must satisfy the useful heat demand:

$$Q_{sup,t} = Q_{total,t} \tag{2.16}$$

where $Q_{sup,t}$ is the useful heat supplied to the household at time step t. This balance relation links the thermal demand model to the specific heating technology adopted.

The thermal demand profile depends on building characteristics, outdoor weather conditions, occupancy patterns, and hot water consumption behavior. In energy system studies, such demand is commonly represented as a time series derived from building simulations or demand datasets [34].

Gas Boiler Model

Gas boilers remain one of the most widely used heating technologies in residential buildings because of their technological maturity, dispatchability, and relatively low upfront cost. A gas boiler produces useful heat through the combustion of natural gas and transfers the thermal output to the building heating and domestic hot water systems.

The useful thermal output of the gas boiler can be expressed as:

$$Q_{boiler,t} = \eta_{boiler}\, E_{gas,t} \tag{2.17}$$

where $Q_{boiler,t}$ is the useful thermal energy delivered by the boiler at time step t (kW), η_{boiler} is the boiler efficiency, and $E_{gas,t}$ is the input energy from natural gas consumption.

The operating power of the boiler may also be limited by its nominal capacity:

$$0 \leq Q_{boiler,t} \leq Q_{boiler}^{max} \tag{2.18}$$

where Q_{boiler}^{max} is the maximum useful heat that can be supplied by the boiler. In simplified residential energy models, the boiler efficiency is often treated as constant, particularly when the focus is on annual or hourly system-level analysis rather than detailed combustion performance [27].

Air-Source Heat Pump Model

Heat pumps are recognized as one of the most efficient low-carbon technologies for residential heating. An air-source heat pump extracts low-temperature heat from ambient air and upgrades it to a useful temperature level through an electrically driven thermodynamic cycle. Since a heat pump transfers heat rather than generating it directly from fuel combustion, it can deliver several units of thermal energy for each unit of electricity consumed [35].

The performance of a heat pump is commonly characterized by its coefficient of performance (COP), defined as the ratio between useful thermal output and electrical input. The useful heat delivered by the heat pump can be expressed as:

$$Q_{HP,t} = COP_t \, P_{HP,t} \tag{2.19}$$

where $Q_{HP,t}$ is the useful thermal output of the heat pump at time step t (kW), $P_{HP,t}$ is the electrical power consumed by the heat pump (kW), and COP_t is the coefficient of performance.

The heat pump thermal output is limited by the installed heating capacity:

$$0 \leq Q_{HP,t} \leq Q_{HP}^{max} \tag{2.20}$$

where Q_{HP}^{max} is the maximum useful thermal output of the ASHP. The COP may be represented as a constant average value or as a time-varying parameter dependent on ambient conditions. In particular, the performance of an air-source heat pump decreases as the outdoor air temperature becomes lower, which is relevant for winter heating applications [35].

When coupled with photovoltaic generation, an ASHP can increase the electrification of residential heating and reduce direct fossil fuel use. For this reason, air-source heat pumps are increasingly considered a key technology in residential decarbonization pathways.

2.6 EV Charging Demand Representation

Electric vehicle (EV) charging is included in this thesis as an electrical end-use component of the residential demand rather than as a detailed transport model. This choice is consistent with the scope of the study, where EV charging is treated as an additional household electricity demand for the archetypes that own an EV. Accordingly, the EV is not modeled through vehicle dynamics or battery-chemistry equations; instead, it is represented as a controllable charging load that must satisfy a required energy demand within an admissible availability window. This formulation is consistent with common home energy management system (HEMS) studies, where EV charging is modeled as an interruptible shiftable load subject to time-window and energy-delivery constraints [32, 33].

For an archetype k, the total household electricity demand including EV charging can be expressed as:

$$P_{load,k,t} = P_{base,k,t} + P_{EV,k,t} \tag{2.21}$$

where $P_{load,k,t}$ is the total electricity demand of archetype k at time step t, $P_{base,k,t}$ is the baseline household electricity demand excluding EV charging, and $P_{EV,k,t}$ is the EV charging power demand. For archetypes without an EV, the charging term is set equal to zero for all time steps.

Because EV charging can be shifted within the periods in which the vehicle is connected at home, it is modeled as an interruptible shiftable load. Its charging power is constrained by both the maximum charger capacity and the vehicle availability at home:

$$0 \leq P_{EV,k,t} \leq w_{k,t}\, P_{EV,k}^{max} \tag{2.22}$$

where $P_{EV,k}^{max}$ is the maximum charging power of the EV for archetype k, and $w_{k,t}$ is a binary availability parameter equal to 1 when the vehicle is available for charging at time step t, and 0 otherwise. In this way, EV charging is automatically prevented outside the admissible charging window.

To ensure mobility-service feasibility, the EV charging schedule must satisfy a required charging energy over the defined EV day. The daily charging requirement is expressed as:

$$E_{k,d}^{req} = \sum_{t \in \mathcal{T}_{EV}(d)} P_{EV,k,t}\, \Delta t \tag{2.23}$$

where $E_{k,d}^{req}$ is the required EV charging energy for archetype k on day d, $\mathcal{T}_{EV}(d)$ denotes the set of time steps belonging to the EV charging day, and Δt is the duration of the time step. In this thesis, an hourly resolution is adopted, therefore $\Delta t = 1$ h.

This representation allows EV charging to be integrated consistently within the residential electricity-demand model while remaining aligned with the scope of the thesis. The EV is therefore modeled as a flexible electrical service demand that interacts with photovoltaic generation, battery storage, and tariff-based scheduling, without introducing an unnecessary level of transport-system detail.

2.7 Summary

This chapter presented the mathematical models of the main components considered in the residential energy system analyzed in this thesis. The modeling framework included photovoltaic generation, battery energy storage, residential electricity demand, and heating systems. For each component, the main physical and operational relations were introduced in order to define its role in the overall residential energy system.

The photovoltaic model was used to represent local renewable electricity generation, while the battery model described the dynamic evolution of stored electrical energy through charging and discharging processes. Residential electricity demand was formulated as the aggregation of household appliance consumption, with a distinction between non-shiftable and shiftable loads. In addition, the thermal demand model and the representations of the gas boiler and air-source heat pump were introduced to characterize the supply of useful heat for space heating and domestic hot water.

These component-level models provide the technical basis for the integrated

optimization model developed in Chapter 3, where the operational scheduling of the system and the associated decision variables and constraints are formulated within a unified mathematical framework.

Chapter 3

Optimization Model

3.1 Introduction

This chapter presents the optimization framework used in this thesis to represent the operation and technology-selection decisions of the residential energy system. While Chapter 2 introduced the mathematical models of the individual components, the present chapter integrates those component-level representations into a unified optimization structure. The purpose of this chapter is therefore to define the decision variables, objective function, and governing constraints that characterize the optimization problem.

The optimization framework is formulated in a general way so that it can represent different levels of system complexity without changing the underlying mathematical logic. In particular, the framework includes an operational electricity-management layer, which captures the coordination between grid electricity, photovoltaic (PV) generation, battery storage, shiftable appliance demand, and electric-vehicle (EV) charging, as well as a thermal technology-selection layer, which represents the economic choice between alternative heating technologies under a fixed thermal service demand. This separation is important because some parts of the thesis address short-term operational scheduling, whereas others focus on annualized technology comparison. The common role of this chapter is to define the optimization structure from which these applications are derived.

The formulation is developed at hourly resolution over the full year in order to capture the temporal interaction between tariff variation, PV generation, battery operation, household demand, and heating requirements. This temporal resolution is consistent with the prepared input data and with the mathematical structure already used in the current thesis for electricity dispatch, battery state-of-charge dynamics, EV charging requirements, and thermal retrofit comparison [28, 26, 32, 33].

The chapter is organized as follows. First, the sets, indices, and parameters used in the mathematical formulation are defined. Second, the decision variables are introduced. Third, the objective function is presented. Finally, the model constraints are grouped into electricity-balance, PV utilization, battery operation, residential load scheduling, EV charging, thermal-demand, and technology-selection constraints.

The final section discusses the main modeling assumptions and the solution approach adopted in this thesis.

3.2 Sets, indices, and parameters

Sets and indices

- $t \in \mathcal{T}$: hourly time index, with $\mathcal{T} = \{1, 2, \ldots, 8760\}$.
- $k \in \mathcal{K}$: apartment archetype index.
- $a \in \mathcal{A}$: appliance class index for shiftable appliances or thermal technology index, depending on context.
- $j \in J_{k,a}$: task index for extracted runs of appliance class a in archetype k.
- $s \in \Omega_{k,a,j}$: feasible start-time index for task j of appliance class a in archetype k.
- $d \in \mathcal{D}$: day index used for EV charging requirements.

Electrical input parameters

For each hour $t \in \mathcal{T}$:

- π_t [€/kWh]: electricity import tariff.
- P_t^{PV} [kW]: available PV generation at the building connection point.

For each archetype $k \in \mathcal{K}$ and hour $t \in \mathcal{T}$:

- $L_{k,t}^{fix}$ [kW]: fixed non-shiftable electrical load.
- $w_{k,t} \in \{0, 1\}$: EV availability mask, equal to 1 when the EV is parked and charging is allowed.
- $E_{k,d}^{req}$ [kWh]: daily EV charging energy requirement.
- $P_k^{EV,\max}$ [kW]: maximum EV charging power.

Shiftable appliance parameters

For each task $j \in J_{k,a}$:

- $\Delta_{k,a,j}$ [h]: task duration.
- $p_{k,a,j,\ell}$ [kW]: fixed power profile of the task at relative hour ℓ, with $\ell = 0, \ldots, \Delta_{k,a,j} - 1$.

The feasible start set $\Omega_{k,a,j}$ is determined by device-specific and comfort-related scheduling rules.

Battery parameters

- E^B [kWh]: battery energy capacity.
- $SOC^{\min}$ [kWh]: minimum allowable state of charge.
- $SOC^{\max}$ [kWh]: maximum allowable state of charge.
- $P^{ch,\max}$ [kW]: maximum battery charging power.
- $P^{dis,\max}$ [kW]: maximum battery discharging power.
- η^{ch} [-]: battery charging efficiency.
- η^{dis} [-]: battery discharging efficiency.

Grid parameter

- $P^{grid,\max}$ [kW]: maximum building grid import power.

Thermal-system parameters

For each archetype a and hour $t \in \mathcal{T}$:

- $Q_{a,t}^{SH}$ [kW$_{th}$]: space-heating demand.
- $Q_{a,t}^{DHW}$ [kW$_{th}$]: domestic hot water demand.
- $Q_{a,t}^{tot}$ [kW$_{th}$]: total thermal demand.
- p_t^{gas} [€/kWh]: gas tariff.
- p_t^{el} [€/kWh]: electricity tariff for ASHP operation.

Technology-specific parameters include:

- η^B [-]: gas boiler efficiency.
- COP^{fix} [-]: fixed coefficient of performance of the ASHP.
- $c^{B,cap}$ [€/kW]: specific investment cost of the gas boiler.
- $c^{HP,cap}$ [€/kW]: specific investment cost of the ASHP.
- $C_a^{B,om}$ [€/year]: annual operation and maintenance cost of the gas boiler.
- $C_a^{HP,om}$ [€/year]: annual operation and maintenance cost of the ASHP.
- r [-]: real discount rate.
- n_B [years]: gas boiler lifetime.
- n_{HP} [years]: ASHP lifetime.

The design thermal capacity for each archetype is defined from the peak thermal demand:

$$P_a^{des} = \max_{t \in \mathcal{T}} Q_{a,t}^{tot} \tag{3.1}$$

To annualize investment costs, the capital recovery factor is defined as:

$$CRF(r, n) = \frac{r(1+r)^n}{(1+r)^n - 1} \tag{3.2}$$

3.3 Decision variables

Electrical dispatch variables

For each hour $t \in \mathcal{T}$, the model includes the following continuous variables:

- $G_t \geq 0$: grid import power [kW],
- $P_t^{PV \to L} \geq 0$: PV power directly supplied to load [kW],
- $P_t^{ch,PV} \geq 0$: battery charging power from PV [kW],
- $P_t^{ch,G} \geq 0$: battery charging power from grid [kW],
- $P_t^{dis} \geq 0$: battery discharge power [kW],
- $P_t^{curt} \geq 0$: curtailed PV power [kW],
- SOC_t: battery state of charge [kWh].

A binary variable is used to enforce non-simultaneous charging and discharging:

$$u_t \in \{0, 1\} \qquad \forall t \in \mathcal{T} \tag{3.3}$$

Shiftable-task variables

For each archetype $k \in \mathcal{K}$, appliance class $a \in \mathcal{A}$, task $j \in J_{k,a}$, and feasible start time $s \in \Omega_{k,a,j}$, the model defines:

$$x_{k,a,j,s} \in \{0, 1\} \tag{3.4}$$

where $x_{k,a,j,s} = 1$ indicates that task j starts at hour s.

EV charging variables

For each relevant archetype $k \in \mathcal{K}$ and hour $t \in \mathcal{T}$:

$$p_{k,t}^{EV} \geq 0 \tag{3.5}$$

denotes EV charging power [kW].

Thermal technology-selection variable

For each archetype a, the thermal technology choice is represented by a binary variable:

$$y_a \in \{0, 1\} \tag{3.6}$$

with

$$y_a = \begin{cases} 0, & \text{if a gas boiler is selected} \\ 1, & \text{if an ASHP is selected} \end{cases}$$

3.4 Objective function

The unified optimization model is formulated as an annual cost-minimization problem. In order to preserve a general master formulation that can later support multiple application cases, the objective function is written as the sum of four cost components:

$$\min Z = F_1 + F_2 + F_3 + F_4 \tag{3.7}$$

where F_1 represents grid-electricity procurement cost, F_2 represents the annualized cost associated with the shared photovoltaic system, F_3 represents the annualized cost associated with the shared battery energy storage system, and F_4 represents the annualized cost of the selected thermal supply technology.

Electricity procurement cost

The first objective term accounts for electricity purchased from the grid over the full optimization horizon:

$$F_1 = \sum_{t \in \mathcal{T}} \pi_t G_t \Delta t \tag{3.8}$$

where π_t [€/kWh] is the time-dependent electricity import tariff, G_t [kW] is the grid import power at hour t, and $\Delta t = 1$ h is the model time step.

Because the formulation is developed at hourly resolution, the above term represents the annual electricity procurement cost associated with operating the building under the considered tariff structure.

Annualized photovoltaic cost

The second objective term represents the annualized cost of the shared PV system:

$$F_2 = C^{PV} \tag{3.9}$$

If PV capacity is treated as exogenously specified, C^{PV} is a fixed annualized accounting term associated with the adopted shared asset. If PV size is treated as

an endogenous investment decision in a future extension of the framework, the same term can be written more generally as

$$F_2 = c^{PV,cap} P^{PV,inst} \, CRF(r, n_{PV}) + C^{PV,om} \tag{3.10}$$

where $c^{PV,cap}$ [€/kW] is the specific PV investment cost, $P^{PV,inst}$ [kW] is the installed PV capacity, n_{PV} is the PV lifetime, and $C^{PV,om}$ [€/year] is the annual PV operation and maintenance cost.

Annualized battery cost

The third objective term represents the annualized cost of the shared BESS:

$$F_3 = C^{BESS} \tag{3.11}$$

If battery size is exogenously specified, C^{BESS} is treated as a fixed annualized accounting term. In the more general case of endogenous battery sizing, the term can be written as

$$F_3 = c^{BESS,cap} E^B \, CRF(r, n_{BESS}) + C^{BESS,om} \tag{3.12}$$

where $c^{BESS,cap}$ [€/kWh] is the specific BESS investment cost, E^B [kWh] is installed battery capacity, n_{BESS} is battery lifetime, and $C^{BESS,om}$ [€/year] is annual battery operation and maintenance cost.

This formulation preserves the generality of the optimization framework while remaining consistent with cases in which PV and battery assets are predefined rather than optimized as investment variables.

Thermal technology cost

The fourth objective term represents the annualized thermal-system cost for each archetype:

$$F_4 = \sum_{a \in \mathcal{A}} \left[(1 - y_a) C_a^B + y_a C_a^{HP} \right] \tag{3.13}$$

where $y_a \in \{0, 1\}$ is the thermal technology-selection variable for archetype a, such that

$$y_a = \begin{cases} 0, & \text{if a gas boiler is selected} \\ 1, & \text{if an ASHP is selected} \end{cases} \tag{3.14}$$

The annual total cost of the boiler option is defined as

$$C_a^B = C_a^{B,cap} + C_a^{B,om} + C_a^{B,fuel} \tag{3.15}$$

and the annual total cost of the ASHP option is defined as

$$C_a^{HP} = C_a^{HP,cap} + C_a^{HP,om} + C_a^{HP,el} \tag{3.16}$$

The annualized capital cost of the boiler is

$$C_a^{B,cap} = c^{B,cap} P_a^{des}\, CRF(r, n_B) \tag{3.17}$$

and the annualized capital cost of the ASHP is

$$C_a^{HP,cap} = c^{HP,cap} P_a^{des}\, CRF(r, n_{HP}) \tag{3.18}$$

where P_a^{des} [kW$_{th}$] is the design thermal capacity for archetype a, determined from the peak hourly thermal demand as

$$P_a^{des} = \max_{t \in \mathcal{T}} Q_{a,t}^{tot} \tag{3.19}$$

The variable operating cost of the gas boiler is

$$C_a^{B,fuel} = \sum_{t \in \mathcal{T}} p_t^{gas} E_{a,t}^{gas} \Delta t \tag{3.20}$$

while the variable electricity cost of the ASHP is

$$C_a^{HP,el} = \sum_{t \in \mathcal{T}} p_t^{el} E_{a,t}^{HP} \Delta t \tag{3.21}$$

The thermal objective term therefore represents the minimum annualized cost of satisfying the given thermal service demand through one of the two alternative technologies.

Expanded form of the master objective

By substituting the component definitions into Eq. (3.7), the objective function can be written in expanded form as

$$\min Z = \sum_{t \in \mathcal{T}} \pi_t G_t \Delta t + C^{PV} + C^{BESS} + \sum_{a \in \mathcal{A}} \Big[(1 - y_a) C_a^{B} + y_a C_a^{HP} \Big] \tag{3.22}$$

or, by further replacing the thermal technology costs,

$$\min Z = \sum_{t \in \mathcal{T}} \pi_t G_t \Delta t + C^{PV} + C^{BESS} + \sum_{a \in \mathcal{A}} \Big[(1 - y_a)(C_a^{B,\mathrm{cap}} + C_a^{B,\mathrm{om}} + C_a^{B,\mathrm{fuel}}) + y_a (C_a^{HP,\mathrm{cap}} + C_a^{HP,\mathrm{om}} + C_a^{HP,\mathrm{el}}) \Big] \tag{3.23}$$

This expanded form is useful because it makes explicit how the different annual cost components contribute to the total objective value.

Interpretation within the thesis structure

The objective function is intentionally written in a general decomposed form. This does not mean that all terms must be active in every application case. Instead, the present chapter defines the full optimization structure, while the scenario chapter activates only the relevant subsets of cost terms, variables, and constraints for each case study. In this way, the mathematical logic remains consistent across the thesis while allowing different scenario configurations to be derived from the same master formulation.

3.5 Constraint set

The unified optimization framework combines hourly electrical dispatch, shared PV and BESS operation, shiftable appliance scheduling, EV charging, and annual thermal technology selection. The full set of constraints is introduced here in general form.

Shiftable appliance task constraints

Selected household appliances are represented as shiftable tasks. Each task $j \in J_{k,a}$ for appliance class a and archetype k has a fixed duration $\Delta_{k,a,j}$ and a fixed power profile

$$\mathbf{p}_{k,a,j} = \left[p_{k,a,j,0}, p_{k,a,j,1}, \ldots, p_{k,a,j,\Delta_{k,a,j}-1}\right].$$

The feasible start set $\Omega_{k,a,j}$ contains all admissible start times such that the full task remains inside the allowed operating interval:

$$\Omega_{k,a,j} = \left\{ s \in \mathcal{T} : h(s+\ell) \in [h_a^{start}, h_a^{end}) \;\; \forall \ell = 0, \ldots, \Delta_{k,a,j} - 1 \right\} \tag{3.24}$$

where $h(t)$ returns the local hour-of-day associated with time index t.

Each task must be scheduled exactly once:

$$\sum_{s \in \Omega_{k,a,j}} x_{k,a,j,s} = 1 \qquad \forall k \in \mathcal{K},\ a \in \mathcal{A},\ j \in J_{k,a} \tag{3.25}$$

To prevent physically impossible overlap of tasks belonging to the same appliance class in the same archetype, the following non-overlap constraint is imposed:

$$\sum_{j \in J_{k,a}} \sum_{s \in \Omega_{k,a,j}} x_{k,a,j,s}\, \mathbb{I}\{t \in [s, s + \Delta_{k,a,j} - 1]\} \leq 1 \qquad \forall k \in \mathcal{K},\ a \in \mathcal{A},\ t \in \mathcal{T} \tag{3.26}$$

where $\mathbb{I}\{\cdot\}$ is an indicator function equal to 1 if the condition is true and 0 otherwise.

The scheduled load of appliance class a for archetype k is reconstructed as

$$P_{k,a,t}^{task} = \sum_{j \in J_{k,a}} \sum_{s \in \Omega_{k,a,j}} x_{k,a,j,s}\, p_{k,a,j,t-s} \qquad \forall k, a, t \tag{3.27}$$

with $p_{k,a,j,t-s} = 0$ whenever $t - s \notin \{0, \ldots, \Delta_{k,a,j} - 1\}$.

The total shiftable load for archetype k is then

$$P_{k,t}^{shift} = \sum_{a \in \mathcal{A}} P_{k,a,t}^{task} \qquad \forall k \in \mathcal{K},\ t \in \mathcal{T} \tag{3.28}$$

This formulation represents load shifting rather than generic same-hour flexibility, because the appliance service is moved from one admissible time interval to another while preserving the task itself.

EV charging constraints

For each archetype $k \in \mathcal{K}$ and hour $t \in \mathcal{T}$, EV charging is constrained by vehicle availability:

$$0 \le p_{k,t}^{EV} \le P_k^{EV,\max} w_{k,t} \qquad \forall k \in \mathcal{K},\ t \in \mathcal{T} \tag{3.29}$$

To satisfy mobility requirements, a daily charging-energy constraint is imposed:

$$\sum_{t \in \mathcal{T}(d)} p_{k,t}^{EV} \Delta t = E_{k,d}^{req} \qquad \forall k \in \mathcal{K},\ d \in \mathcal{D} \tag{3.30}$$

where $\mathcal{T}(d)$ denotes the set of hours belonging to day d.

PV balance and battery operation

The PV generation available at the building meter is treated as an exogenous input time series. At each hour, the available PV power is split among direct consumption, battery charging, and curtailment:

$$P_t^{PV \to L} + P_t^{ch,PV} + P_t^{curt} = P_t^{PV} \qquad \forall t \in \mathcal{T} \tag{3.31}$$

Battery state of charge evolves according to

$$SOC_t = SOC_{t-1} + \eta^{ch} \left(P_t^{ch,PV} + P_t^{ch,G} \right) \Delta t - \frac{1}{\eta^{dis}} P_t^{dis} \Delta t \qquad \forall t \in \mathcal{T} \tag{3.32}$$

subject to

$$SOC^{\min} \le SOC_t \le SOC^{\max} \qquad \forall t \in \mathcal{T} \tag{3.33}$$

A cyclic boundary condition is imposed to avoid end-of-horizon bias:

$$SOC_0 = SOC_T \tag{3.34}$$

A binary variable enforces non-simultaneous charging and discharging:

$$P_t^{ch,PV} + P_t^{ch,G} \leq P^{ch,\max} u_t \qquad \forall t \in \mathcal{T} \tag{3.35}$$

$$P_t^{dis} \leq P^{dis,\max}(1 - u_t) \qquad \forall t \in \mathcal{T} \tag{3.36}$$

If grid charging is enabled in a given application, it is additionally constrained by

$$P_t^{ch,G} \leq G_t \qquad \forall t \in \mathcal{T} \tag{3.37}$$

Building electrical balance

The total electrical load of archetype k is defined as

$$L_{k,t} = L_{k,t}^{fix} + P_{k,t}^{shift} + p_{k,t}^{EV} \qquad \forall k \in \mathcal{K},\ t \in \mathcal{T} \tag{3.38}$$

and the aggregated building load is

$$L_t^{bld} = \sum_{k \in \mathcal{K}} L_{k,t} \qquad \forall t \in \mathcal{T} \tag{3.39}$$

The hourly building-level electrical balance is

$$P_t^{PV \to L} + P_t^{dis} + G_t = L_t^{bld} + P_t^{ch,G} \qquad \forall t \in \mathcal{T} \tag{3.40}$$

No export is allowed. This is enforced by construction because the formulation does not include any export variable; PV surplus that cannot be directly used or stored must be curtailed.

To avoid unrealistic concentration of load in low-price hours, grid import can also be restricted by the building connection limit:

$$G_t \leq P^{grid,\max} \qquad \forall t \in \mathcal{T} \tag{3.41}$$

Thermal energy-conversion constraints

For each archetype a and hour t, the total thermal demand is

$$Q_{a,t}^{tot} = Q_{a,t}^{SH} + Q_{a,t}^{DHW} \qquad \forall a \in \mathcal{A},\ t \in \mathcal{T} \tag{3.42}$$

If the thermal demand is supplied by a gas boiler with efficiency η^B, the hourly gas input is

$$E_{a,t}^{gas} = \frac{Q_{a,t}^{tot}}{\eta^B} \qquad \forall a \in \mathcal{A},\ t \in \mathcal{T} \tag{3.43}$$

If the thermal demand is supplied by an ASHP with fixed coefficient of performance COP^{fix}, the hourly electricity input is

$$E_{a,t}^{HP} = \frac{Q_{a,t}^{tot}}{COP^{fix}} \qquad \forall a \in \mathcal{A},\ t \in \mathcal{T} \tag{3.44}$$

The technology-selection variable y_a determines which annual cost is active in the objective function for each archetype.

Archetype-level allocation of shared electricity cost

When household-level billing is required for a shared behind-the-meter system, the building-level grid import can be allocated to archetypes using a time-resolved proportional sharing rule based on instantaneous demand share.

Let the aggregated load of archetype k be $L_{k,t}$ and the building load be L_t^{bld}. The demand share of archetype k is defined as

$$\alpha_{k,t} = \begin{cases} \dfrac{L_{k,t}}{L_t^{bld}}, & \text{if } L_t^{bld} > 0 \\ 0, & \text{if } L_t^{bld} = 0 \end{cases} \tag{3.45}$$

The allocated grid import is then

$$G_{k,t} = \alpha_{k,t} G_t \qquad \forall k \in \mathcal{K},\ t \in \mathcal{T} \tag{3.46}$$

and the corresponding allocated annual electricity cost is

$$C_{el,k}^{alloc} = \sum_{t \in \mathcal{T}} \pi_t G_{k,t} \Delta t \tag{3.47}$$

This allocation is budget-balanced by construction:

$$\sum_{k \in \mathcal{K}} C_{el,k}^{alloc} = \sum_{t \in \mathcal{T}} \pi_t G_t \Delta t \tag{3.48}$$

Model class and solution approach

The unified optimization framework contains both continuous and binary decision variables. When only continuous dispatch variables are active, the formulation reduces to a linear program (LP). When binary variables for battery charge/discharge exclusivity, appliance start times, or thermal technology selection are activated, the formulation becomes a mixed-integer linear program (MILP).

The model is implemented in Python using the PuLP optimization package and solved with the CBC solver. This formulation is suitable for coupled residential energy-management problems because it can represent both continuous energy flows and discrete scheduling decisions while preserving linearity in the objective function and the constraint set.

Main modeling assumptions

The optimization framework is based on the following assumptions:

- The model is solved at hourly resolution over one representative year.
- PV generation is treated as an exogenous time series at the building connection point.
- Grid export is not allowed; unused PV surplus is curtailed.
- Battery operation is modeled with fixed charging and discharging efficiencies.
- Battery state of charge is subject to lower and upper bounds and a cyclic end-of-year condition.
- Shiftable household appliances are modeled as non-interruptible tasks with fixed duration and fixed power profile.
- EV charging is modeled as interruptible, subject to an availability mask and a daily energy requirement.
- Comfort-related restrictions are represented through feasible operating windows rather than monetary penalty terms.
- Thermal-system comparison is based on deterministic sizing from peak hourly demand and annualized cost evaluation using the capital recovery factor.
- The general framework is broader than any single scenario and is intentionally written so that specific cases can later be created by activating only the relevant subsets of variables, constraints, and cost terms.

3.6 Summary

This chapter presented the unified optimization framework adopted in the thesis. The formulation combines hourly electricity procurement, shared PV and BESS operation, shiftable appliance scheduling, EV charging, and annual thermal technology selection within a common cost-minimization structure. The resulting objective terms, decision variables, and constraints provide the mathematical basis for the scenario ladder introduced in Chapter 4, where specific study cases are generated by activating only the relevant subsets of the general model.

Chapter 4

Scenario Ladder

4.1 Introduction

This chapter defines the scenario ladder used to apply the optimization framework introduced in Chapter 3. While Chapter 3 established the general mathematical structure of the model, including the decision variables, objective-function components, and governing constraints, the present chapter specifies how that common framework is translated into a set of progressively structured case studies. The purpose of the scenario ladder is to evaluate the contribution of different system configurations and service-management options in a transparent and comparable way.

The scenarios are arranged in increasing order of system complexity. Scenario 0 defines the baseline reference case without shared energy assets or service scheduling. Scenario 1 introduces a shared photovoltaic system and a shared battery energy storage system at building level. Scenario 2 extends this configuration by allowing service scheduling for selected residential end-uses within admissible operating windows. Scenario 3 addresses the thermal retrofit dimension by comparing alternative thermal supply options under the same general methodological framework. In this way, each scenario isolates a specific decision layer while remaining linked to the same overall modeling logic.

This chapter does not re-derive the optimization model. Instead, each scenario is defined by identifying the relevant subsets of objective terms, decision variables, and constraints already introduced in Chapter 3. This structure enables consistent scenario comparison while avoiding unnecessary repetition of the mathematical formulation. It also makes clear which technological features and operational decisions are activated in each case.

For each scenario, the chapter describes its definition and scope, the main model components that are activated, and the output indicators used later in the results chapter. The reported indicators include economic, operational, and energy-related quantities such as annual total cost, grid electricity procurement, PV utilization, battery operation, and the comparative performance of thermal retrofit options. The scenario ladder therefore provides the bridge between the general optimization

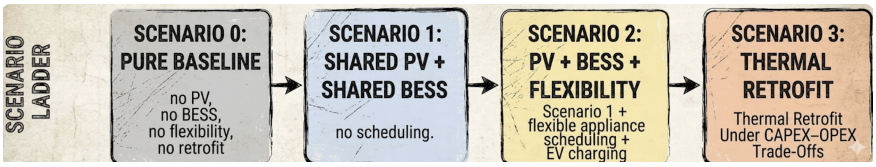

Figure 4.1: Scenario Ladder.

framework and the simulation results presented in Chapter 5.

4.2 Scenario 0: Baseline

Definition and scope

Scenario 0 defines the baseline reference case used throughout the scenario ladder. It represents the conventional operating condition of the residential building before the introduction of shared photovoltaic generation, shared battery storage, service scheduling, or thermal retrofit measures. The scenario is formulated as a *service-preserving* reference case, meaning that the same household energy services are maintained as in the subsequent scenarios while no additional technology or scheduling intervention is activated. This framing is important because it avoids misleading comparisons in which apparent savings would result from a reduction in delivered service rather than from an actual improvement in system operation or technology choice [36].

In practical terms, Scenario 0 assumes that household electricity demand is supplied entirely by grid imports under the applicable hourly time-of-use tariff, while gas demand is met through the existing gas-based configuration without any technology substitution. Residential electricity demand is treated as fully exogenous, including the baseline contribution of electric-vehicle charging where relevant to the archetype definition, and no scheduling decision is introduced for appliances or mobility demand. The baseline therefore reflects the original demand pattern of the building, without any coordination mechanism or behind-the-meter optimization.

The role of Scenario 0 is not to improve performance, but to establish the benchmark against which all subsequent scenarios are evaluated. Since Scenario 1 introduces shared PV and shared BESS, Scenario 2 adds service scheduling, and Scenario 3 examines thermal retrofit options, a clearly defined reference case is required in order to quantify the incremental value of each intervention. Scenario 0 therefore provides the benchmark annual procurement cost, the benchmark pattern of grid dependence, and the benchmark gas-cost level from which all scenario comparisons are derived.

Activated model components

Within the general framework introduced in Chapter 3, Scenario 0 activates only the baseline procurement logic associated with exogenous demand and imported energy. The electricity-cost component of the objective function is retained in its baseline form, as introduced in Section 3.4, while the PV-related term, the BESS-related

term, the service-scheduling variables, and the thermal technology-selection variable remain inactive in this scenario.

At the operational level, the building electricity balance is reduced to the baseline condition in which total building demand is supplied directly by the grid. The building load is obtained by aggregating the archetype-level electricity demands across the building, consistently with the aggregation logic presented in Chapter 3. In the same way, baseline gas demand is aggregated from the archetype level and valued using the gas-price representation adopted in the thesis.

Because no shared PV system is active in Scenario 0, the PV utilization, battery charging, battery discharging, and curtailment relations defined in Chapter 3 are not activated. Likewise, the task-based service-scheduling constraints and EV scheduling variables introduced in the general model remain inactive, since all residential loads are treated as fixed demand inputs in the baseline case. The thermal technology-selection formulation is also excluded, because Scenario 0 preserves the existing thermal-service supply structure rather than comparing alternative technologies.

Cost representation

All costs reported in Scenario 0 are procurement-based operating costs. For electricity, the annual cost is determined by the interaction between the fixed hourly building-demand profile and the hourly time-of-use electricity tariff. For gas, the annual baseline cost is determined by the exogenous gas-demand series and the gas-price signal adopted in the thesis. The total annual baseline cost is therefore obtained by combining the electricity and gas procurement components under the original service-demand pattern of the building.

The baseline annual cost is evaluated first at archetype level and then aggregated to building level. For each archetype, the total annual baseline procurement cost is defined as the sum of the corresponding electricity and gas costs. The building-level baseline cost is then obtained by weighting each archetype cost by the number of apartments represented by that archetype and summing over all archetypes. This two-level representation is useful because it allows the thesis to report both detailed archetype-level benchmark values and the overall building-level reference cost used in the scenario comparison.

This cost representation is consistent with the role of Scenario 0, which is to provide a transparent benchmark rather than to optimize technology choice or dispatch decisions. It also ensures that the later scenarios can be compared against a common cost reference under the same underlying service requirements. In this way, any reduction in annual cost observed in subsequent scenarios can be interpreted as the result of improved system operation, local generation, storage support, service scheduling, or thermal technology substitution, rather than as a change in the level of service delivered to residents [37, 36].

Purpose within the scenario ladder

Scenario 0 is the anchor of the scenario ladder. Its purpose is to answer the following baseline question:

> What is the annual energy procurement cost and grid-dependence level of the building when no shared energy asset, scheduling strategy, or retrofit measure is introduced?

This question is fundamental because all later scenarios are interpreted relative to this reference. Scenario 1 measures the value of shared PV and shared BESS relative to Scenario 0. Scenario 2 evaluates the additional benefit of service scheduling on top of the shared-energy configuration. Scenario 3 compares thermal retrofit options against the original thermal-service supply condition. Without Scenario 0, these incremental effects could not be interpreted rigorously.

From a methodological point of view, Scenario 0 confirms that the thesis is not comparing unrelated system configurations, but rather a sequence of progressively enhanced cases that preserve the same service demand while changing only the technological or operational layer under study. This is the essential logic of the scenario ladder.

Reported outputs

The outputs reported for Scenario 0 are selected to provide the benchmark indicators needed for the comparative analysis in Chapter 5. These outputs are reported both at archetype level and at building level, depending on the purpose of the analysis.

The main reported outputs are:

- annual electricity consumption and annual electricity procurement cost,
- electricity consumption and electricity cost split by time-of-use band,
- annual gas consumption and annual gas procurement cost,
- total annual baseline procurement cost,
- hourly aggregated building-load profile,
- hourly grid-import profile under baseline conditions.

These indicators establish the economic and operational benchmark against which the later scenarios are evaluated. In particular, they provide the reference values used to quantify cost savings, reductions in grid import, changes in time-of-use exposure, and the comparative value of technology and scheduling interventions.

Consistency checks

Although Scenario 0 does not involve optimization decisions, it still requires careful consistency checks because all subsequent scenarios depend on its validity as a reference case. For this reason, the baseline implementation is verified through a set of automated checks applied to the hourly input series and annual aggregations.

The verification procedure includes checks for time-series completeness over the 8,760-hour annual horizon, timestamp alignment across electricity, gas, EV, and tariff data, non-negativity of demand and price series, consistency of time-of-use energy and cost aggregation, and enforcement of zero EV demand for archetypes without EV ownership. These checks are essential because an incorrect baseline would distort every later comparison in the scenario ladder.

The full verification log is archived together with the Scenario 0 outputs so that the reported benchmark values remain traceable and reproducible.

4.3 Scenario 1: Shared PV + Shared BESS

Definition and scope

Scenario 1 extends the baseline configuration by introducing a shared photovoltaic (PV) system and a shared battery energy storage system (BESS) at the building level. The scenario preserves the same residential service demands defined in Scenario 0 and does not introduce service scheduling or thermal technology replacement. In other words, the demand side remains exogenous, while the technological configuration is enhanced through shared local generation and shared storage operating behind the building meter.

The purpose of Scenario 1 is to isolate the operational value of shared PV generation and shared battery dispatch under the same time-of-use electricity tariff used in the baseline case. Relative to Scenario 0, this scenario allows the building to partially replace grid electricity purchases with local renewable generation and to shift part of the electricity use over time through battery charging and discharging. The comparison with Scenario 0 therefore reveals the economic and operational effect of introducing shared energy assets while keeping the same underlying service demand.

Scenario 1 also includes a PV-only interpretation within the broader comparison logic. This is useful because it allows the contribution of PV self-consumption to be distinguished from the additional value created by battery operation. In this way, the scenario can separately highlight the role of direct PV utilization and the incremental benefit of shared storage under the same demand and tariff conditions.

Activated model components

Within the general optimization framework introduced in Chapter 3, Scenario 1 activates the electricity procurement component together with the shared PV and

shared BESS operational relations. In particular, the building-level electricity balance, the PV allocation logic, the battery state-of-charge dynamics, the charging and discharging limits, and the no-export condition introduced in Chapter 3 are active in this scenario.

The demand side remains unchanged with respect to Scenario 0. Residential electricity demand is still treated as an exogenous hourly input, including the baseline EV charging contribution where relevant to the archetype definition. No service-scheduling variables are activated, and the timing of appliance use and mobility demand is therefore not modified in this scenario. Likewise, the thermal technology-selection formulation remains inactive, because Scenario 1 focuses exclusively on the operational interaction between shared PV generation, shared battery storage, and grid electricity procurement.

At the operational level, the scenario represents a shared behind-the-meter energy system. PV generation can be used directly to supply building load, can be stored in the battery, or can be curtailed when it cannot be absorbed locally. Battery operation is used to reduce grid procurement cost by shifting electricity over time within the battery operational limits. Since the scenario assumes no remunerated export to the grid, surplus PV that is not directly consumed or stored is treated as curtailed energy. This no-export assumption is consistent with the shared behind-the-meter configuration adopted in the thesis and allows the value of self-consumption and storage to be evaluated directly.

Cost representation

In Scenario 1, the electricity-related cost is no longer determined only by the fixed interaction between building demand and the tariff. Instead, it depends on the optimized coordination between grid imports, shared PV generation, and shared battery dispatch. The electricity procurement term introduced in Chapter 3 remains the operative part of the annual electricity-cost calculation, while the reduction in grid imports relative to Scenario 0 captures the economic value of local generation and storage operation.

Gas demand and gas procurement cost remain unchanged with respect to the baseline case, because Scenario 1 does not modify the thermal-service supply structure. The total annual cost reported for Scenario 1 therefore combines the optimized electricity procurement cost with the unchanged gas-cost component inherited from Scenario 0. This ensures that the economic comparison remains fully service-preserving and directly comparable with the baseline case.

Because the shared PV and BESS are modeled at building level, electricity-cost reporting is performed first for the shared building system and then, where needed, allocated back to archetypes using the proportional allocation rule defined in Chapter 3. This makes it possible to report both building-level results and archetype-level allocated electricity costs without changing the physical operation of the shared system.

Purpose within the scenario ladder

Scenario 1 is the first technological extension of the baseline. Its purpose is to answer the following question:

> How much economic and operational value can be obtained by introducing shared PV generation and shared battery storage, while keeping residential demand fixed and preserving the same service level?

This question is central to the thesis because it isolates the effect of shared energy assets before any demand-side scheduling is introduced. In the scenario ladder, Scenario 1 therefore plays an intermediate role between the purely exogenous baseline and the more advanced coordinated-demand case of Scenario 2. It establishes whether shared renewable generation and storage alone are sufficient to reduce annual electricity procurement cost and grid dependence, and it provides the reference point against which the additional value of service scheduling can later be measured.

More specifically, Scenario 1 allows the thesis to evaluate three distinct effects: the reduction in grid electricity imports due to direct PV self-consumption, the reduction in procurement cost due to battery-assisted temporal shifting, and the extent to which shared storage reduces PV curtailment by absorbing local surplus generation. These effects cannot be isolated in the baseline case and must therefore be introduced at this stage of the scenario ladder.

Reported outputs

The outputs reported for Scenario 1 are selected to compare the building performance under shared PV and shared BESS operation against the benchmark values established in Scenario 0. These outputs are reported primarily at building level, with archetype-level allocated electricity costs included where relevant for comparative discussion.

The main reported outputs are:

- annual electricity procurement cost,
- annual total cost including the unchanged gas-cost component,
- annual grid electricity import,
- annual PV generation,
- annual PV utilization and annual PV curtailment,
- battery charging and discharging energy,
- battery state-of-charge evolution and related operating statistics,
- archetype-level allocated electricity cost under the shared-system accounting rule.

These indicators are sufficient to characterize the operational and economic role of the shared PV and BESS configuration. In particular, they allow the thesis to quantify electricity-cost savings relative to Scenario 0, the reduction in grid dependence achieved through self-consumption and storage dispatch, and the extent to which shared storage improves the local utilization of PV generation.

Consistency checks

Because Scenario 1 introduces optimized energy flows between the grid, the shared PV system, and the shared battery, additional consistency checks are required beyond those already applied in Scenario 0. For this reason, the implementation is verified through automated checks covering both time-series integrity and operational feasibility.

The verification procedure includes checks for time-series completeness and timestamp alignment, physical plausibility of the PV generation series, hourly power-balance consistency, battery state-of-charge feasibility, enforcement of charge/discharge exclusivity, no-export compliance, and budget balance of the archetype-level electricity-cost allocation. These checks are essential because an apparently favorable result could otherwise arise from an invalid dispatch pattern, inconsistent time alignment, or an incorrect accounting allocation.

The full verification log is archived together with the Scenario 1 outputs so that the shared PV and BESS results remain traceable, reproducible, and directly comparable with the baseline case.

4.4 Scenario 2: Shared PV + Shared BESS + Service Scheduling

Definition and scope

Scenario 2 extends the shared PV and shared BESS configuration of Scenario 1 by introducing service scheduling for selected residential end-uses. In this scenario, the building is still equipped with a shared photovoltaic system and a shared battery energy storage system operating behind the building meter, but part of the demand side is no longer treated as fully fixed. Instead, selected services are allowed to shift in time within admissible operating windows, while the same underlying service requirements are preserved.

The service-scheduling layer includes two main classes of controllable end-uses. The first consists of non-interruptible appliance tasks, such as dishwasher, washing machine, and dryer cycles, whose operation can be shifted in time but must remain continuous once started. The second consists of EV charging, which is modeled as a flexible electrical service subject to availability windows and daily energy requirements. In both cases, the objective is not to reduce the amount of service delivered, but to improve the timing of service provision so that electricity procurement cost can be reduced under the same tariff and technology configuration.

Scenario 2 therefore represents the most advanced operational case in the scenario ladder. Compared with Scenario 1, it adds a demand-side decision layer on top of the shared PV and BESS dispatch problem. This makes it possible to evaluate whether service scheduling provides additional value beyond the value already obtained from local generation and storage alone.

Activated model components

Within the general optimization framework introduced in Chapter 3, Scenario 2 activates the electricity procurement term together with the shared PV and shared BESS operational model already used in Scenario 1. In addition, it activates the service-scheduling variables and constraints associated with shiftable appliance tasks and EV charging. These include the feasible-start-set formulation for non-interruptible appliance cycles, the service-conservation requirement that each task must still be completed, the non-overlap rule for appliance operation, the EV charging bounds linked to vehicle availability, the daily EV energy requirement, and the building-level grid-cap condition.

The physical shared-energy model remains the same as in Scenario 1. PV generation can be used directly to supply load, can charge the battery, or can be curtailed if it cannot be absorbed locally. The battery can shift electricity over time within the state-of-charge and power constraints defined in Chapter 3. What changes in Scenario 2 is that part of the electrical demand is now endogenously scheduled rather than fully exogenous. This means that the total building-load profile becomes the result of both technology operation and service timing decisions.

Comfort and feasibility are preserved explicitly. Non-interruptible appliance tasks are not allowed to run outside their admissible windows, and noisy appliance cycles are excluded from inappropriate hours through hard scheduling constraints rather than through monetary penalty terms. In the same way, EV charging is only permitted when the vehicle is available and must still satisfy the required daily charging energy. This ensures that any improvement in cost is achieved through better coordination rather than through an implicit reduction of service quality.

The thermal technology-selection formulation remains inactive in Scenario 2, since the scenario is still focused on operational electricity management rather than thermal retrofit choice.

Cost representation

As in Scenario 1, the main economic driver of Scenario 2 is the annual electricity procurement cost under the hourly time-of-use tariff. However, unlike Scenario 1, the cost now depends not only on shared PV utilization and battery dispatch, but also on the optimized timing of selected residential services. The procurement cost is therefore reduced, where possible, through the joint coordination of local generation, storage operation, and service scheduling.

Gas demand and gas procurement cost remain unchanged with respect to Scenarios 0 and 1, because Scenario 2 does not alter the thermal-service supply configuration. The total annual cost reported for this scenario therefore combines the optimized electricity procurement cost with the same gas-cost component used in the previous scenarios.

From a comparative point of view, the main economic role of Scenario 2 is to identify the incremental value of service scheduling beyond the shared PV and BESS case. For this reason, Scenario 2 is interpreted primarily relative to Scenario 1. The difference between the two scenarios isolates the value of scheduling appliance tasks and EV charging under the same shared-energy configuration and the same service requirements.

Purpose within the scenario ladder

Scenario 2 is the final operational step of the scenario ladder. Its purpose is to answer the following question:

> How much additional economic and operational value can be obtained when selected residential services are scheduled optimally on top of a shared PV and shared BESS configuration?

This question is important because shared generation and shared storage do not fully determine building performance. Even with the same assets, the timing of electricity use strongly affects the degree of PV self-consumption, the effectiveness of battery operation, the exposure to high-price tariff periods, and the maximum grid import required by the building. Scenario 2 therefore tests whether coordinated service timing can unlock value that remains unavailable when all demand is fixed.

Within the overall scenario ladder, Scenario 2 represents the highest level of operational coordination. Scenario 0 provides the benchmark, Scenario 1 introduces shared assets, and Scenario 2 adds active demand-side management while preserving the same service demand. The comparison between Scenario 1 and Scenario 2 is therefore particularly important, because it reveals whether service scheduling contributes meaningful additional value beyond what has already been achieved by shared PV and shared battery storage alone.

Reported outputs

The outputs reported for Scenario 2 are selected to evaluate both the economic effect of service scheduling and its operational interaction with the shared PV and shared BESS system. These outputs are reported primarily at building level, with selected archetype-level indicators included where relevant.

The main reported outputs are:

- annual electricity procurement cost,
- annual total cost including the unchanged gas-cost component,

- annual grid electricity import,
- annual PV utilization and annual PV curtailment,
- battery charging and discharging energy,
- battery state-of-charge evolution and related operating statistics,
- scheduled appliance-load profile,
- EV charging profile under optimized scheduling,
- comparison between baseline and scheduled timing of selected service tasks,
- incremental cost reduction and incremental grid-import reduction relative to Scenario 1.

These indicators make it possible to assess whether service scheduling improves the economic use of shared PV and shared BESS, whether it reduces grid dependence further than Scenario 1, and how the timing of flexible services changes under the optimized solution. They also support the interpretation of whether the value of scheduling is substantial or only marginal once shared generation and storage are already present.

4.5 Scenario 3: Thermal Retrofit Evaluation

Definition and scope

Scenario 3 shifts the focus of the scenario ladder from operational electricity management to thermal technology choice. While Scenarios 0–2 examine the economic and operational effects of shared PV generation, shared battery storage, and service scheduling on the electricity side, Scenario 3 evaluates the thermal retrofit dimension of the residential building by comparing alternative technologies for supplying the same thermal services.

The scenario is formulated as a service-preserving thermal retrofit evaluation. The required thermal services, namely space heating and domestic hot water, remain unchanged, and only the technology used to deliver those services is allowed to vary. In this way, the comparison does not depend on any reduction in comfort or service quality, but only on the economic and energetic consequences of adopting one thermal supply option instead of another.

In the present thesis, the comparison is carried out between two mutually exclusive alternatives: a gas boiler configuration and an air-source heat pump configuration. The purpose is therefore not to optimize short-term scheduling across multiple interacting thermal devices, but to determine which technology provides the required thermal service at lower annualized cost under the same hourly service-demand profile.

Activated model components

Within the general optimization framework introduced in Chapter 3, Scenario 3 activates the thermal-demand representation, the thermal energy-conversion relations, and the thermal technology-selection component of the objective function. In particular, the hourly thermal service demand, the conversion of that service demand into gas or electricity input depending on the selected technology, the annualized capital-cost formulation, and the binary technology-choice logic are active in this scenario.

By contrast, the operational electricity-management components used in Scenarios 1 and 2 are not active here. Shared PV generation, shared BESS operation, service scheduling, EV scheduling, and the associated electrical dispatch constraints are excluded from Scenario 3, because the purpose of this scenario is not to study coordinated building-level electricity operation, but to isolate the techno-economic effect of changing the thermal conversion technology while preserving the same service demand.

The thermal comparison is therefore narrower in scope than the preceding scenarios, but this narrower scope is deliberate. It allows the thesis to evaluate the retrofit decision directly, without mixing it with the short-term operational effects of electrical flexibility or shared storage dispatch.

Scenario 3 logic overview

Before discussing the economic results, Figure 4.2 summarizes the logic of Scenario 3. The hourly thermal-service demand is obtained by aggregating space-heating and domestic-hot-water demand for each archetype. The same useful thermal demand is then supplied through two alternative technology paths: a condensing gas boiler and an air-source heat pump (ASHP). For each option, the corresponding annualized total cost is calculated by combining operating cost, annualized capital cost, and annual operation and maintenance cost. The optimization then selects the lower-cost technology for each archetype through a binary technology-choice variable.

Cost representation

The economic comparison in Scenario 3 is based on annualized total cost. For each candidate thermal technology, the total annual cost combines two components: the annualized capital-related cost and the annual operating cost required to supply the given hourly thermal demand. This makes it possible to compare a lower-investment but fuel-intensive technology with a higher-investment but potentially more efficient electrified alternative on a common yearly basis.

For the gas boiler option, the annual cost includes the annualized investment cost, annual operation and maintenance cost, and fuel-related operating cost. For the air-source heat pump option, the annual cost includes the annualized investment cost, annual operation and maintenance cost, and electricity-related operating cost associated with supplying the same thermal demand. The cost comparison is therefore

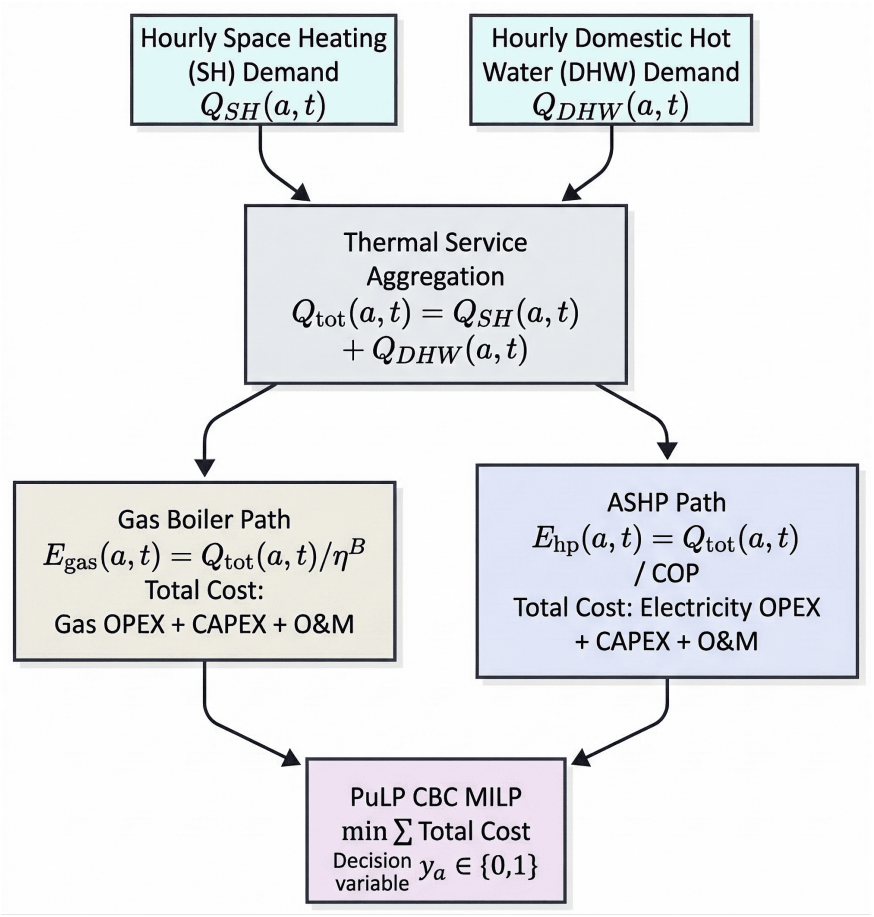

Figure 4.2: Logic of the Scenario 3 thermal retrofit comparison.

technology-neutral at the level of delivered service, because both options are evaluated against the same hourly thermal-service requirement.

From the perspective of the scenario ladder, Scenario 3 differs from the previous scenarios in one important way. In Scenarios 0–2, the main economic question concerns operational procurement cost under different electrical system configurations. In Scenario 3, the central question is instead whether the thermal retrofit option changes the annualized cost of delivering the same thermal services. The cost representation is therefore investment-aware and service-preserving rather than purely dispatch-based.

Purpose within the scenario ladder

Scenario 3 provides the thermal retrofit dimension of the thesis. Its purpose is to answer the following question:

> Which thermal supply technology provides the required space-heating and domestic-hot-water services at lower annualized cost under the same thermal demand conditions?

This question is important because the energy transition of residential buildings is not determined only by electricity-side flexibility and self-consumption. The choice of thermal conversion technology can also reshape both the economic structure of

energy use and the building's dependence on different energy carriers. A gas boiler preserves the conventional fuel-based thermal supply pathway, whereas an air-source heat pump shifts the thermal service to electricity and may therefore support broader electrification of the residential energy system.

Within the scenario ladder, Scenario 3 complements the operational scenarios rather than extending them directly. Scenario 0 defines the reference condition, Scenario 1 adds shared PV and shared BESS, Scenario 2 adds service scheduling, and Scenario 3 evaluates whether a thermal retrofit measure changes the annualized cost of supplying the same thermal services. It therefore introduces the technology-substitution perspective that completes the thesis beyond short-term dispatch optimization alone.

Reported outputs

The outputs reported for Scenario 3 are selected to compare the techno-economic performance of the two thermal supply options under the same service-demand conditions. These outputs are reported at archetype level and, where relevant, aggregated to building level.

The main reported outputs are:

- annualized total cost of the gas boiler option,
- annualized total cost of the air-source heat pump option,
- annualized capital cost of each thermal technology,
- annual operating cost of each thermal technology,
- design thermal capacity required for each archetype,
- annual gas input associated with the boiler option,
- annual electricity input associated with the air-source heat pump option,
- selected thermal technology under the annualized cost comparison,
- cost difference between the two thermal retrofit options.

These indicators are sufficient to evaluate whether the retrofit option is economically favorable and to explain the relative role of investment cost and operating cost in the final decision. They also provide the basis for discussing whether the thermal retrofit changes the cost structure of the building by replacing direct gas use with electrified thermal-service provision.

4.6 Summary

This chapter translated the general optimization framework introduced in Chapter 3 into a structured scenario ladder for the residential case study. Scenario 0 established

the baseline reference condition, Scenario 1 introduced shared photovoltaic generation and shared battery energy storage, Scenario 2 extended the shared-energy configuration by adding service scheduling for selected residential end-uses, and Scenario 3 addressed the thermal retrofit dimension through the comparison of alternative thermal supply options. Together, these scenarios define the analytical pathway used in the thesis to evaluate the economic and operational value of progressively more advanced residential energy interventions.

The chapter also clarified the role of each scenario within the overall logic of the thesis. Scenario 0 provides the reference condition against which all later cases are assessed. Scenario 1 isolates the operational value of shared local generation and storage. Scenario 2 examines whether coordinated service timing can unlock additional value beyond that obtained from shared assets alone. Scenario 3 complements the operational analysis by evaluating the annualized cost implications of thermal technology substitution under preserved service demand. In this way, the scenario ladder links the general model formulation of Chapter 3 with the comparative results presented in Chapter 5.

Chapter 5

Data Preparation and Study Case

5.1 Introduction

This chapter defines the study case and prepares the input datasets required to implement the optimization framework developed in Chapter 3. The analysis is applied to the Italian context, with the case-study building located in Turin, and all exogenous inputs are organized for a full-year simulation horizon corresponding to the year 2025. The choice of Italy is motivated by the relevance of the national residential energy transition, the increasing role of distributed photovoltaic generation, the continued importance of natural gas in residential thermal services, and the need to evaluate how electrification, storage, and load shifting can reduce annual energy procurement costs at building level.

The purpose of this chapter is to move from the abstract methodological framework to an operational case-study model that can be simulated consistently across all scenarios. While Chapter 3 established the mathematical structure of the optimization problem, including the decision variables, objective-function terms, and constraints, the present chapter explains how the case-study system is constructed from real and synthetic datasets, how the required energy services are represented, and how all time series are converted into model-ready hourly inputs. This step is essential because the credibility of the optimization results depends not only on the model formulation itself, but also on the coherence, completeness, and consistency of the data used to drive it.

The model year is set to 2025. This provides a unified temporal reference for all electricity, thermal, photovoltaic, meteorological, and tariff inputs used in the study. Adopting a single model year avoids inconsistencies that would arise if datasets from different years were combined without harmonization. It also allows the thesis to represent the operation of the residential energy system over a complete annual cycle, including winter heating demand, summer solar availability, weekday–weekend variations, and seasonal changes in household electricity use. In this sense, the 2025 horizon is not treated as a speculative forecast of future household behaviour, but as

the common simulation year used to assemble a coherent case-study database.

To make the modelling process transparent, the data are organized into clearly defined sectors. These include the case-study definition and modelling boundary, the household electricity-demand dataset, the thermal-demand dataset, the mobility-related demand dataset where applicable, the electricity and gas price datasets, and the weather and photovoltaic dataset. This structure strengthens traceability and makes it possible to verify each input stream separately before it enters the optimization framework.

A further objective of this chapter is to preserve the service-based logic of the thesis. The model does not begin from technology dispatch alone; it begins from the services required by the occupants, such as appliance electricity use, domestic hot water, space heating, and, where relevant, electric mobility. These services are then translated into hourly demand profiles and linked to the technologies considered in the scenario analysis. This modelling choice is important because it ensures that scenario comparisons remain fair: any reduction in annual cost or grid dependence must arise from better coordination, improved technology choice, or more efficient conversion, rather than from an artificial reduction in the level of service delivered to residents.

The chapter also addresses household heterogeneity within the building. Rather than assuming that all apartments follow the same demand pattern, the case study represents multiple household archetypes with different occupancy structures and activity profiles. This improves realism at building level because diversity in electricity and thermal demand affects the timing of peaks, the coincidence between load and PV generation, and the operational value of shared storage and load shifting. A homogeneous load assumption would be simpler to implement, but it would reduce the representativeness of the case study and could bias the estimated value of shared energy resources.

Another central task of this chapter is data harmonization. Since the model integrates several data streams, all inputs must be converted to a common hourly resolution and aligned over the same annual index of 8,760 hours. This includes timestamp parsing, unit conversion, treatment of missing values, aggregation logic, and consistency checks between service demand and final energy demand. Without this harmonization step, the optimization outputs could appear mathematically valid while remaining physically misleading. For that reason, this chapter is not a merely descriptive part of the thesis; it is the stage at which the case-study database is made robust and reproducible.

In summary, this chapter answers a practical question that is central to the thesis: how is the Italian 2025 residential case study built for simulation? The answer is developed by defining the modelling boundary, specifying the household structure, organizing the required input datasets, and transforming each data stream into a form suitable for the optimization framework. Once these steps are completed, the case study is ready for the scenario simulations reported in the following chapters.

5.2 Study case

This section defines the residential case study and the system boundary adopted in the thesis. All exogenous input datasets are prepared at hourly resolution for the year 2025 and aligned to a common annual time index of 8,760 time steps.

Building description and system boundary

The case study is a multi-family residential building located in Turin, Italy, and composed of 20 dwelling units. For weather and solar-resource retrieval, as well as for reproducibility of the photovoltaic input data, the building location is represented by the reference coordinates (45.044°N, 7.639°E), corresponding approximately to an elevation of 257 m above sea level. The exact street address is omitted for privacy, while the geographical representation is sufficient for the construction of meteorological and PV-related inputs.

The electrical boundary is defined at the building point of connection to the public distribution grid. The building exchanges electricity with the grid through hourly imports and, where applicable, hourly exports. The on-site electrical resources considered in the thesis are a rooftop photovoltaic (PV) system and, in the scenarios that include storage, a shared building-level battery energy storage system (BESS). Both resources are connected on the building electrical side and interact with the aggregated building demand.

At the apartment level, electricity demand is represented by one hourly load profile per dwelling. The total building electricity demand is obtained by aggregation:

$$P_{\text{load,building}}(t) = \sum_{i=1}^{20} P_{\text{load},i}(t), \tag{5.1}$$

where $P_{\text{load},i}(t)$ denotes the hourly electricity demand of apartment i.

Thermal services are represented through the final energy carrier that supplies them. When thermal demand is met by gas-based technologies, the relevant carrier is natural gas. When thermal demand is supplied by electrified technologies, such as a heat pump, the corresponding carrier is electricity. This dual-carrier representation is consistent with the service-based modelling logic adopted throughout the thesis and allows the optimization framework to compare different technology options under a common cost-minimization structure.

Household archetypes and allocation to dwelling units

Household heterogeneity is represented through four archetypes designed to capture distinct occupancy and activity patterns within the same residential building:

- working couple,
- couple with one child,

Figure 5.1: Geographical location of the case-study building (Turin, Italy).

Table 5.1: Case-study overview (boundary-level parameters).

Item	Value
Location	Turin, Italy
Reference coordinates	45.044°N, 7.639°E (approx. 257 m a.s.l.)
Building size	20 dwellings (multi-family residential)
Time resolution	1 hour
Time horizon (2025)	8,760 hours
Modeled energy carriers	Electricity and natural gas
Household archetypes	Four types
On-site resources (scenarios)	Rooftop PV and optional shared BESS

- couple with three children,
- retired couple.

Hourly electricity demand profiles for each archetype are generated synthetically using the LoadProfileGenerator (LPG) tool, which produces behaviour-consistent residential load profiles by simulating occupant presence, daily activities, and appliance usage patterns [38, 39]. In the absence of measured smart-meter data for the specific case-study building, this approach provides a reproducible basis for obtaining full-year hourly demand traces while preserving controlled household heterogeneity.

In the base building composition, the 20 apartments are distributed evenly across the four archetypes, with five dwellings assigned to each type. This allocation is used consistently across the thesis to construct apartment-level demand profiles and to obtain the final building-level load by aggregation. The chosen composition is not intended to describe a unique real building with exact census accuracy; rather, it is adopted as a representative and tractable case-study structure that preserves diversity in demand timing and magnitude.

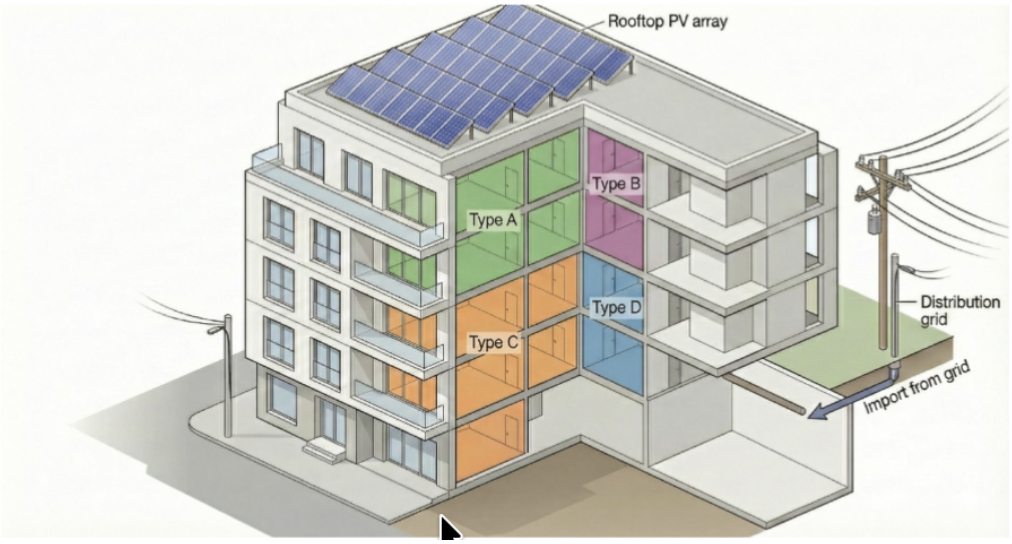

Figure 5.2: Case-study building and modelling boundary.

Service-to-device mapping and modelled energy carriers

The demand side of the case study is framed in terms of *energy services* rather than a fixed list of appliances. The services required by occupants are mapped to device operations only to the extent necessary for constructing hourly input datasets and representing shiftable demand within the optimization. This service-based framing is consistent with the structure introduced in Chapter 3 and ensures that the modelling focus remains on satisfying household needs rather than on prescribing device-level behaviour independently of service delivery.

Within this framework, household electricity demand is divided into two broad categories:

- **Non-shiftable services**, corresponding to baseline uses that must be satisfied in the hour in which they occur, such as lighting, plug loads, refrigeration, standby loads, and other non-deferrable household uses.
- **Shiftable service-tasks**, corresponding to device operations that can be shifted within admissible time windows without changing the final service delivered to occupants, such as laundry cycles, dishwasher cycles, and electric-vehicle charging.

Two final energy carriers are considered in the thesis:

- **Electricity**, used for household electrical services and for electrified technologies, including EV charging and heat-pump-based thermal supply.
- **Natural gas**, used for thermal services when these are supplied by gas-based technologies, such as a boiler or an instantaneous gas water heater.

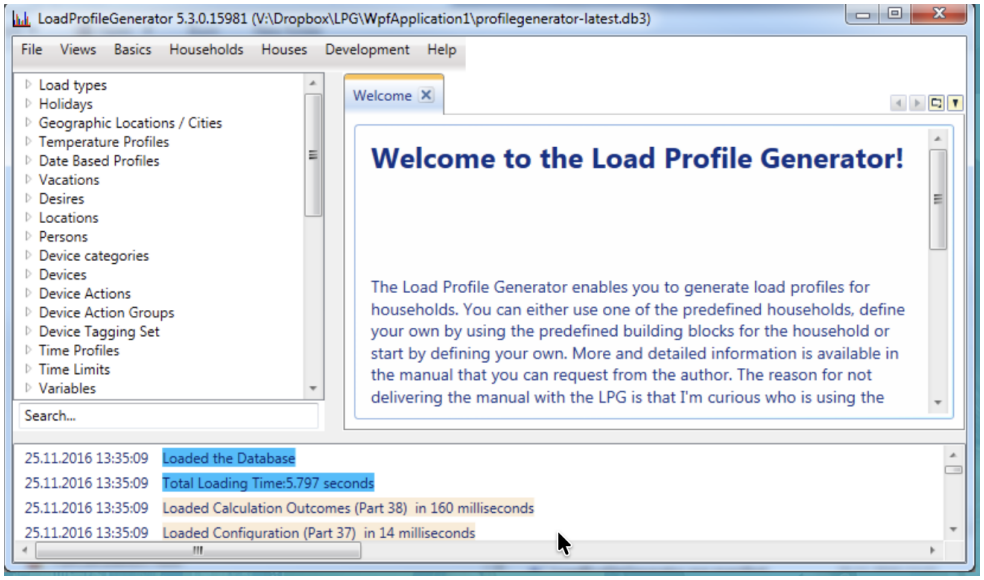

Figure 5.3: LoadProfileGenerator interface for residential load profile generation.

The building geometry used to assess photovoltaic potential is extracted from an OpenStreetMap-derived building polygon dataset. The required attribution is therefore given according to OpenStreetMap guidance [40]. The detailed PV sizing assumptions and the PV time-series generation workflow are discussed later in Section 5.8 and reported in full detail in Appendix A.

5.3 Service-demand profiles

This section describes how the hourly service-demand inputs are constructed for the optimization model, starting from archetype electricity profiles generated in LoadProfileGenerator (LPG) and extending them into a tractable service representation suitable for hourly optimization. The goal is to maintain a service-based interpretation of residential demand while ensuring compatibility with the mixed-integer linear programming framework.

Electricity demand profiles from LoadProfileGenerator

For each household archetype, an hourly electricity-demand profile is generated using LoadProfileGenerator, an agent-based simulation environment that produces synthetic but behaviour-consistent residential load traces by modelling occupancy, activities, and appliance usage [38, 39]. LPG is used in this thesis to obtain full-year hourly profiles with controlled heterogeneity in the absence of measured smart-meter data for the case-study building.

The simulation settings adopted for the archetype runs are summarized in Table 5.2. The outputs are exported at household level in hourly energy units and subsequently aligned to the common annual index used in the optimization model. A concise summary of the LPG configuration is also reported in Appendix B.2.

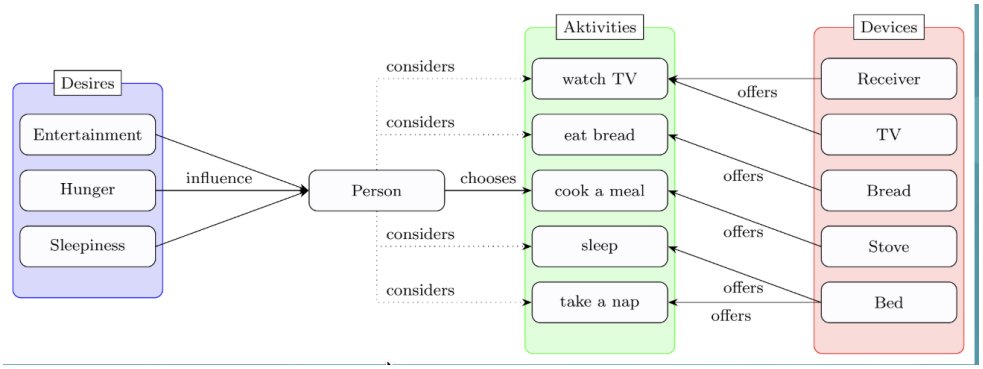

Figure 5.4: LPG logic from occupant activities to device-level electricity demand.

Table 5.2: LPG configuration for archetype electricity-demand profiles.

Item	Value
Simulation period	01/01/2025–31/12/2025 (full year)
Time resolution	3600 s (1 hour)
Energy carrier	Electricity; Gas
Geographical setting	European residential context (weekday/holiday calendars)
Archetypes	Working couple; 1-child; 3-children; retired couple
Outputs exported	Hourly household electricity demand (kWh/h)

The use of synthetic demand data requires an explicit modelling assumption: the generated profiles are not direct measurements of a specific building, but representative hourly traces consistent with the activity patterns of the selected household types. Their role in the thesis is therefore not to reconstruct the exact consumption history of one real building, but to provide a coherent and reproducible demand dataset suitable for case-study simulation and scenario comparison.

From household electricity demand to service categories

The optimization model does not operate on a fixed list of independent appliances. Instead, it uses a service-based demand representation in which each apartment demand profile is decomposed into a non-shiftable baseline component and a set of shiftable service-tasks. For each apartment i, the hourly electricity demand is represented as:

$$P_{\text{load},i}(t) = P_i^{\text{NS}}(t) + \sum_{k \in \mathcal{K}_i} P_{i,k}(t), \tag{5.2}$$

where $P_i^{\text{NS}}(t)$ is the non-shiftable demand component, $\mathcal{K}_i$ is the set of shiftable task categories considered for apartment i, and $P_{i,k}(t)$ denotes the contribution associated with task type k.

This decomposition provides two advantages. First, it preserves the physical interpretation of the total load as the sum of essential and shiftable uses. Second, it creates a direct interface with the optimization model, in which selected tasks can be rescheduled within admissible time windows while the non-shiftable demand remains

fixed.

A compact summary of the main service categories and their carrier mapping is reported in Table 5.3.

Table 5.3: Service categories, example end uses, and main energy carriers.

Category	**Example end uses**	**Main carrier**
Non-shiftable services	lighting, plug loads, baseline household use	Electricity
Shiftable services (non-interruptible)	laundry cycle, dishwasher cycle	Electricity
Shiftable services (interruptible)	EV charging	Electricity
Thermal services	space heating, domestic hot water, cooking	Electricity and/or gas

Task types, shiftability, and admissible time windows

To keep the demand representation consistent with common residential scheduling formulations, two classes of shiftable service-tasks are distinguished:

- **Non-interruptible tasks**: once started, they must run continuously until completion. Typical examples include washing-machine, dryer, and dishwasher cycles.
- **Interruptible tasks**: the required energy can be delivered in multiple time intervals within an admissible availability window, provided that the total required energy is satisfied by the deadline. In this thesis, EV charging is modelled in this category.

This distinction is important because it preserves behavioural and technical plausibility. Non-interruptible household cycles are not artificially split into discontinuous fragments, while EV charging retains the scheduling flexibility that is more consistent with vehicle parking patterns. In the optimization model introduced previously, non-interruptible tasks are therefore associated with consecutive-hour operation constraints, whereas interruptible EV charging is represented through availability and energy-by-deadline constraints.

The time windows assigned to shiftable tasks are selected to preserve realism without overcomplicating the behavioural assumptions. For noisy household tasks, the admissible window is restricted to daytime and evening hours in order to avoid unrealistic night-time operation. For EV charging, the admissible periods reflect residential availability, especially overnight and evening presence at home. In this way, the case study captures *load shifting* rather than unrestricted scheduling.

Post-processing and alignment of LPG outputs

All exported LPG time series are post-processed to ensure compatibility with the hourly optimization horizon. The post-processing includes:

- **timestamp parsing and alignment**, by converting the exported series to a continuous hourly index;
- **completeness checks**, by verifying the presence of 8,760 hourly samples for the non-leap simulation year;
- **sanity checks**, including non-negativity and absence of missing values;
- **feature extraction**, including hour-of-day, day-of-week, and month indicators used in later analyses and tariff mapping.

This post-processing stage is necessary because the optimization model requires all exogenous inputs to share a common time base. The household load dataset is therefore not used directly after generation; it is first converted into a clean, aligned, and model-ready hourly input.

5.4 Thermal-demand dataset

This section describes how thermal demand is represented in the case study. The thermal side of the building is treated at the level of required services, with final energy consumption depending on the technology used to supply those services.

Thermal services and carrier mapping

The main thermal services considered in the thesis are space heating and domestic hot water (DHW). These services are represented independently of the technology that provides them. This means that the service requirement exists first, and the conversion from final energy to useful heat depends on the selected supply technology.

When space heating or DHW is supplied by gas-based technologies, the relevant input carrier is natural gas. When the same service is supplied by electrified technologies, such as an air-source heat pump, the relevant carrier becomes electricity. This separation is important because it allows the optimization to compare technology options on a common service basis instead of embedding technology assumptions directly into the demand definition.

Domestic hot water volume to gas consumption

The available DHW dataset provides an hourly profile of hot-water production volume, expressed in litres per hour. To estimate gas consumption associated with this service, the DHW volume is first converted into useful thermal energy delivered to the water, then mapped to fuel input energy through an efficiency assumption, and finally converted into standard cubic metres of natural gas.

Let V_t denote the DHW volume produced in hour t in litres, T_{hot} the hot-water setpoint temperature, and $T_{\text{cold},t}$ the cold mains-water temperature at hour t. Assuming water density $\rho_w \approx 1\,\text{kg/L}$ and specific heat capacity $c_w \approx 4.186\,\text{kJ/(kg K)}$, the useful thermal energy delivered to the water is:

$$Q_t^{\text{DHW}} = \rho_w V_t c_w \left(T_{\text{hot}} - T_{\text{cold},t}\right). \tag{5.3}$$

Expressed in kWh, this becomes:

$$Q_t^{\text{DHW}}\,[\text{kWh}] \approx V_t\,[\text{L}] \cdot \left(T_{\text{hot}} - T_{\text{cold},t}\right)\,[^\circ\text{C}] \cdot 0.001163, \tag{5.4}$$

where the factor 0.001163 converts the product of mass and temperature difference into kWh. This formulation is consistent with the DHW energy-balance approach used in UNI/TS 11300-2 [41].

When direct measurements of mains-water temperature are unavailable, the cold-water temperature is represented through a smooth seasonal approximation. Following a sinusoidal approach commonly used in the literature [42], the Turin mains-water temperature is approximated as:

$$T_{\text{cold}}(d) = T_{\text{mean}} + A \sin\!\left(\frac{2\pi(d - d_0)}{365}\right), \tag{5.5}$$

where d is the day of the year. In this thesis, the parameters are set to $T_{\text{mean}} = 14^\circ\text{C}$, $A = 6^\circ\text{C}$, and $d_0 = 134$, yielding a realistic annual range of approximately 8°C in winter and 20°C in late summer for Turin.

Let η_{DHW} denote the efficiency of the instantaneous gas water heater. The hourly fuel input energy is:

$$E_t^{\text{gas}}\,[\text{kWh}] = \frac{Q_t^{\text{DHW}}}{\eta_{\text{DHW}}}, \tag{5.6}$$

with a conservative baseline value of $\eta_{\text{DHW}} = 0.82$ adopted in the absence of detailed equipment specifications. This is consistent with literature-based values for non-condensing gas tankless heaters, while condensing units can reach higher efficiencies [43].

To convert the gas input energy into standard cubic metres (Smc), the higher heating value (PCS) is used:

$$V_t^{\text{gas}}\,[\text{Smc}] = \frac{E_t^{\text{gas}}}{\text{PCS}}, \tag{5.7}$$

where PCS depends on gas composition and is reported on Italian bills [44]. A representative value of PCS $\approx 10.7\,\text{kWh/Smc}$ is adopted for the present case study.

The annual DHW energy and associated gas consumption are finally obtained by hourly aggregation:

$$Q_{\text{year}}^{\text{DHW}} = \sum_t Q_t^{\text{DHW}}, \qquad E_{\text{year}}^{\text{gas}} = \sum_t E_t^{\text{gas}}, \qquad V_{\text{year}}^{\text{gas}} = \sum_t V_t^{\text{gas}}. \tag{5.8}$$

This conversion procedure ensures that the DHW dataset is expressed in a form directly compatible with the energy-balance and cost calculations used later in the optimization.

5.5 Mobility-related demand dataset

Electric-vehicle charging demand is included in the case study when the relevant scenario design requires mobility-related electricity consumption. In the present thesis, EV charging is associated with the working-couple and couple-with-one-child archetypes, reflecting the more likely presence of commuting-related mobility demand in these household types.

The EV dataset is represented as an hourly electricity-demand series and aligned to the same 2025 annual time base used for the other exogenous inputs. Where the source trace is archetype-based, it is scaled to building level by multiplication with the number of apartments assigned to the relevant household type.

Daily energy requirement and overnight alignment

Because EV charging windows often cross midnight, a simple calendar-day aggregation from 00:00 to 24:00 can artificially split a single overnight charging opportunity into two separate daily constraints. To avoid this inconsistency, daily EV energy requirements are defined over an offset 24-hour interval aligned with the end of the overnight charging window. For each day d:

$$E_{k,d}^{\mathrm{req}} = \sum_{t \in \mathcal{T}_{EV}(d)} E_{k,t}^{EV,\mathrm{base}}, \tag{5.9}$$

where $\mathcal{T}_{EV}(d)$ denotes the shifted daily interval used to define the EV charging requirement for archetype k.

This approach improves consistency between daily energy requirements and the admissible charging window, especially in cases where evening arrival and morning departure define the physically relevant availability period.

Presence-based availability windows

An hourly availability mask $w_{k,t}$ is constructed to represent the periods during which the EV is assumed to be connected and available for charging. The mask is based on the type of day (weekday, Saturday, Sunday, or national holiday) and follows the basic principle that the vehicle is mainly available at home during evening and overnight hours on weekdays, with wider availability on weekends and holidays.

This modelling approach is consistent with common residential EV-charging formulations, where availability is linked to arrival and departure patterns rather than treated as unrestricted charging over the full day [45, 46]. In the optimization model, the EV charging power is therefore constrained both by availability and by the daily energy requirement defined above.

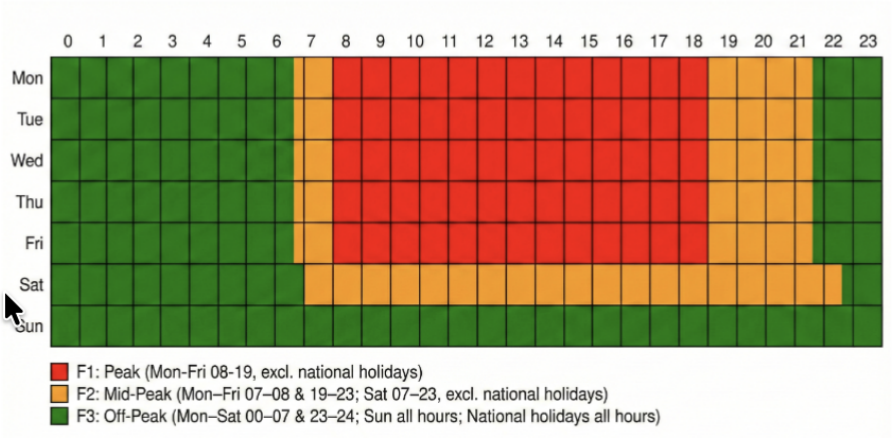

Figure 5.5: Italian TOU bands used in this thesis (F1/F2/F3).

Table 5.4: Fixed TOU import prices used in the 2025 tariff dataset.

Band	Price [/kWh]
F1	0.129865
F2	0.120466
F3	0.106513

5.6 Electricity tariff dataset

An hourly time-of-use (TOU) electricity-price dataset is constructed for the full year 2025. Each hour is assigned to one of the Italian retail time bands $b(t) \in \{\mathrm{F1}, \mathrm{F2}, \mathrm{F3}\}$ according to the ARERA definition of electricity time bands [47].

The band mapping adopted in this thesis is the following:

- **F1 (peak)**: Monday–Friday, 08:00–19:00, excluding national holidays;
- **F2 (mid-peak)**: Monday–Friday, 07:00–08:00 and 19:00–23:00, and Saturday 07:00–23:00, excluding national holidays;
- **F3 (off-peak)**: Monday–Saturday 00:00–07:00 and 23:00–24:00, and all hours on Sundays and national holidays.

Once the band assignment is defined, a fixed import price is assigned to each hour through:

$$p^{\mathrm{imp}}(t) = p_{\mathrm{F1}} \cdot \mathbb{I}[b(t) = \mathrm{F1}] + p_{\mathrm{F2}} \cdot \mathbb{I}[b(t) = \mathrm{F2}] + p_{\mathrm{F3}} \cdot \mathbb{I}[b(t) = \mathrm{F3}], \tag{5.10}$$

where p_{F1}, p_{F2}, and p_{F3} are the band-specific electricity prices taken from the supplier offer adopted in this thesis. The resulting values are reported in Table 5.4.

National holidays are assigned to F3 for all 24 hours, consistently with the ARERA band definition [47]. The list of Italian national holidays used in the 2025 mapping is reported in Appendix B.1.

The resulting hourly price series constitutes the electricity procurement signal used in the optimization model. Its role is not merely descriptive: it directly influences the value of battery dispatch, the timing of shiftable electricity uses, and the comparative economics of electrified thermal technologies.

5.7 Natural gas price dataset

Household gas tariffs do not provide an hourly price signal equivalent to electricity TOU pricing. In Italy, ARERA describes the gas bill as the sum of several macro-components, including commodity sales, transport and meter management, system charges, and taxes or VAT [48]. For this reason, two complementary gas-price representations are used in the thesis:

1. a **commodity-only gas price signal**, suitable for marginal cost comparisons and technology sensitivity;
2. a **full retail gas cost representation**, suitable for annual bill-consistent comparisons.

Commodity-only gas price

As a public and reproducible proxy for the commodity component of natural gas, the monthly ARERA index $CMEM_m$ is adopted, expressed in /Smc and published for the "Servizio di tutela della vulnerabilità" [49]. Since the optimization is hourly, the monthly value is mapped to all hours within the corresponding calendar month:

$$p_h^{\text{gas}} = CMEM_{m(h)} \qquad [/\text{Smc}], \tag{5.11}$$

where $m(h)$ denotes the month associated with hour h.

When conversion to /kWh is required, the ARERA calorific reference of $PCS = 0.038520$ GJ/Smc is used:

$$1 \text{ Smc} \approx 0.038520 \text{ GJ/Smc} \times 277.78 \text{ kWh/GJ} \approx 10.7 \text{ kWh}, \tag{5.12}$$

and therefore:

$$p_m^{\text{gas}} [/\text{kWh}] = \frac{CMEM_m [/\text{Smc}]}{10.7}. \tag{5.13}$$

The monthly values used for 2025 are summarized in Table 5.5.

Full retail gas cost representation

A single full-retail price obtained by dividing total bill cost by total gas consumption is not appropriate as a marginal price because the gas bill includes fixed charges that depend on contract structure rather than on hourly usage [48, 50]. For this reason, the full-retail representation separates variable and fixed components.

Table 5.5: ARERA $CMEM_m$ values used in this thesis for 2025 (/Smc).

Month (2025)	$CMEM_m$ [/Smc]
Jan	0.533576
Feb	0.566178
Mar	0.455069
Apr	0.402365
May	0.403010
Jun	0.418839
Jul	0.392478
Aug	0.380886
Sep	0.373358
Oct	0.353669
Nov	0.348704
Dec	0.327985

Let $\mathcal{V}$ denote the set of variable gas-bill components expressed per Smc. The monthly variable retail price is:

$$p_m^{\text{gas,var}} = \sum_{j \in \mathcal{V}} c_{j,m} \qquad [/\text{Smc}], \tag{5.14}$$

and the corresponding hourly marginal gas cost is:

$$p_h^{\text{gas,var}} = p_{m(h)}^{\text{gas,var}}. \tag{5.15}$$

The fixed components of the bill are aggregated at annual level:

$$C_{\text{year}}^{\text{gas,fix}} = \sum_{j \in \mathcal{F}} C_j, \tag{5.16}$$

where $\mathcal{F}$ includes all fixed quota items. This annual term is relevant for overall techno-economic comparison, especially when evaluating whether a gas connection is retained or replaced by full electrification, but it does not influence intra-day scheduling decisions.

This two-level representation of gas cost allows the thesis to distinguish correctly between hourly marginal operation and annual bill-consistent comparison.

5.8 Rooftop PV sizing and generation

This section defines the photovoltaic input used in the case study. It includes both the rooftop PV sizing envelope, which sets the technical upper bound for PV capacity, and the hourly PV generation series used as the exogenous renewable input for the 2025 simulation year.

Table 5.6: PV sizing envelope used in the thesis.

Parameter	Value
Roof area from OSM polygon	$A_{\text{roof}} = 642\ \text{m}^2$
Maximum PV modules	$N_{\text{mod}}^{\max} = 127$
PV size decision	$0 \leq S_{\text{PV}} \leq S_{\text{PV}}^{\max}$
Upper bound definition	$S_{\text{PV}}^{\max} = N_{\text{mod}}^{\max} P_{\text{mod}}$

Rooftop PV sizing envelope

The roof geometry is derived from an OpenStreetMap-based building polygon and used to define the physically feasible rooftop area available for PV deployment. The usable roof area is:

$$A_{\text{roof}} = 642\ \text{m}^2. \tag{5.17}$$

Based on the available footprint and a conservative layout assumption that accounts for spacing, walkways, setbacks, and partial shading, the maximum feasible number of PV modules is:

$$N_{\text{mod}}^{\max} = 127. \tag{5.18}$$

Let P_{mod} denote the rated module power in kWp per module. The PV peak-power upper bound is then:

$$S_{\text{PV}}^{\max} = N_{\text{mod}}^{\max} P_{\text{mod}}. \tag{5.19}$$

In the scenarios that include PV sizing as a decision, the installed PV capacity is constrained by:

$$0 \leq S_{\text{PV}} \leq S_{\text{PV}}^{\max}. \tag{5.20}$$

Hourly PV time-series generation for 2025

The optimization requires an hourly PV generation input for the reference year 2025. Since the adopted workflow aims to remain consistent with PVGIS-based PV performance estimation while also producing a full 2025 series, the PV input is generated using a PVGIS-consistent data-preparation procedure. Historical PVGIS hourly outputs are used as the target variable, while hourly meteorological and radiation predictors are retrieved from Open-Meteo [51, 52, 53].

The regression model is based on gradient-boosted decision trees (XGBoost), which are used to approximate the mapping between the meteorological feature vector $\mathbf{x}_t$ and the corresponding PV specific yield:

$$\hat{g}_t = f_\theta(\mathbf{x}_t), \tag{5.21}$$

where $\hat{g}_t$ is the predicted hourly specific PV generation and f_θ is the trained regression function [54]. The predictor set includes radiation variables such as global, direct, and diffuse irradiance, as well as cloud cover, temperature, wind speed, and calendar information [53].

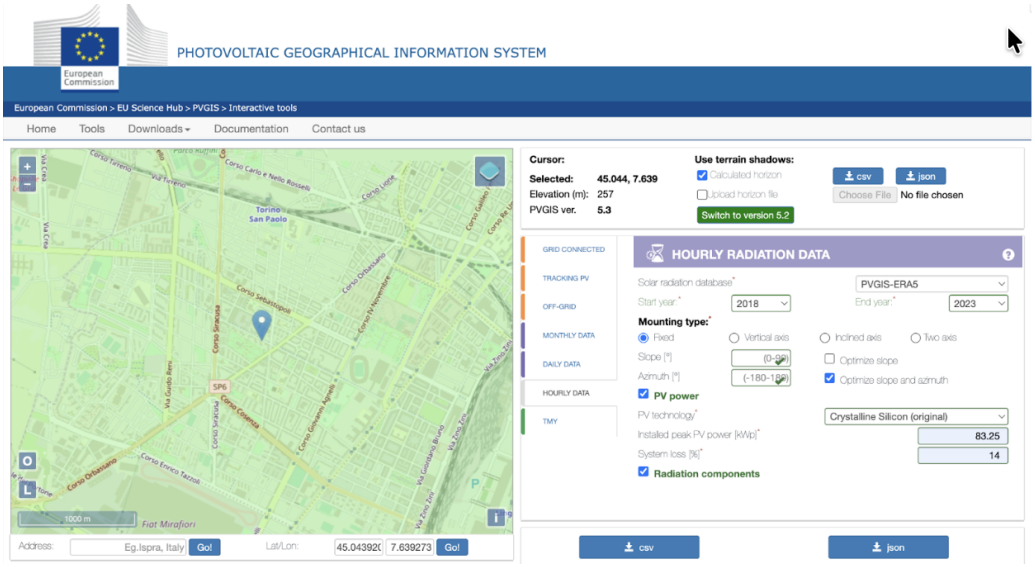

Figure 5.6: PVGIS interactive tool interface.

To ensure physical plausibility, the predicted values are post-processed to enforce non-negativity:

$$\hat{g}_t \leftarrow \max(0, \hat{g}_t), \tag{5.22}$$

and zero output is imposed during night-time hours where irradiance indicates the absence of solar input.

The hourly PV generation time series used in the model is then obtained by linear scaling with the installed PV capacity:

$$PV_t = \hat{g}_t \, S_{\mathrm{PV}}. \tag{5.23}$$

This expression is used consistently across the thesis. In fixed-size cases, S_{PV} is a parameter. In investment cases, it is a decision variable bounded by Eqs. (5.20)–(5.19). In both cases, the resulting series PV_t acts as the exogenous hourly renewable-generation input for the electricity balance.

The detailed technical workflow for PV sizing, time-series generation, and validation is intentionally reported in Appendix A, so that the main chapter can remain focused on the role of the PV series as an input dataset for the optimization rather than on low-level implementation details.

Data inventory, harmonization, and consistency checks

This section summarizes the main input datasets used in the case study and describes the harmonization steps required before these data can be introduced into the optimization model.

Table 5.7: Main datasets used in the case study and their role in the model.

Dataset	Resolution / year	Units	Purpose
Archetype electricity traces	hourly, 2025	kWh per hour	Household electricity demand by archetype; basis for apartment-level and building-level load construction.
DHW profile / thermal-service input	hourly, aligned to 2025	L/h or converted kWh/h	Basis for domestic hot water service and associated fuel-consumption estimation.
EV charging trace	hourly, aligned to 2025	kWh	Mobility-related electricity demand where applicable.
PV series (2025)	hourly, 2025	kWh/h or kW equivalent	Exogenous PV generation input PV_t.
TOU electricity tariff	hourly, 2025	/kWh	Exogenous electricity import price series $p^{\mathrm{imp}}(t)$.
Gas price dataset	monthly mapped to hourly, 2025	/Smc or /kWh	Marginal and annual gas-cost representation.

Data inventory

Table 5.7 summarizes the principal datasets used in the thesis and their role in the case-study database.

Common time base and alignment

A single reference time index is defined for the year 2025, covering 8,760 hourly steps. All input series are reindexed or constructed to match this reference. This is a critical modelling choice because it prevents subtle misalignment errors in which demand, generation, and prices would otherwise be evaluated on different timestamps.

The harmonization procedure includes:

1. parsing and standardizing timestamps,
2. sorting and deduplicating observations where needed,
3. converting all inputs to hourly resolution,
4. shifting or aligning source data to the 2025 simulation year where required,
5. reindexing all series to the common hourly DateTimeIndex,
6. checking units and applying any necessary conversions,
7. filling only physically justified missing values.

The use of a common annual time base is particularly important for coupling tariff signals with demand and PV availability. Since the optimization is driven by hourly costs and hourly balances, even minor timestamp inconsistencies could lead to misleading results.

Consistency checks

Before entering the optimization framework, all datasets are subjected to a set of consistency checks. These include:

- **8760-hour completeness**: all aligned series must contain exactly 8,760 hourly values.
- **Index consistency**: all time series must share the same hourly index.
- **Unit consistency**: electricity, gas, and thermal-service inputs must be expressed in units compatible with the mathematical formulation of the model.
- **Non-negativity**: demand and generation series must not contain physically meaningless negative values unless explicitly defined by sign convention.
- **Aggregation consistency**: apartment-level and archetype-level inputs must aggregate correctly to building level.
- **Carrier consistency**: electricity and gas datasets must correspond to the carrier mapping defined by the selected service-supply technology.

For PV data, an additional sanity check is required on the reported unit, since some sources report PV output in watts rather than kilowatts. In such cases, the series must be converted before entering the optimization. For thermal data, the consistency of DHW conversion from volume to useful energy and then to gas consumption must also be verified. These checks ensure that the resulting simulation database is not only complete, but also physically meaningful.

5.9 Summary

This chapter established the case-study foundation of the thesis by defining the Italian residential building, fixing the modelling boundary, and preparing the hourly input datasets for the 2025 simulation year. The required information was organized into distinct sectors, including building definition, electricity demand, thermal demand, mobility-related demand, electricity and gas prices, and weather/PV inputs. Particular attention was given to household heterogeneity, service-based demand representation, rooftop PV sizing, and the conversion of all datasets to a common hourly time base.

By harmonizing these inputs and preserving a consistent service-oriented interpretation of residential energy use, the chapter created the reproducible database required for the optimization framework introduced in Chapter 3. The case study is therefore fully specified at this stage and can be used in the following chapters for the scenario simulations and comparative analysis.

Chapter 6

Simulation Results and Discussion

6.1 Introduction

This chapter presents the simulation results of the scenario ladder introduced in Chapter 4, using the optimization framework developed in Chapter 3 and the case-study data prepared in Chapter 5 The aim of the chapter is not only to report the numerical outputs of the simulations, but also to interpret their meaning in economic and operational terms. For this reason, the discussion focuses on how each scenario modifies the cost structure, energy flows, and service-management behavior of the studied residential building.

The results are organized according to the progressive logic of the scenario ladder. Scenario 0 defines the baseline benchmark and represents the reference condition with no shared photovoltaic generation, no shared battery energy storage system, no shiftable-service scheduling, and no thermal retrofit action. Scenario 1 introduces shared PV and shared BESS in order to evaluate the first technological improvement beyond the baseline. Scenario 2 extends this configuration by adding the scheduling of selected shiftable residential services, thus assessing the incremental value of operational coordination. Scenario 3 addresses the thermal retrofit dimension separately by comparing alternative thermal supply technologies under the same building demand conditions. This structure allows the contribution of each intervention to be assessed step by step, instead of combining all decision layers into a single opaque result.

A key objective of this chapter is therefore comparative interpretation. The baseline scenario provides the annual reference cost and the original dependence on external energy carriers. The following scenarios are then evaluated in relation to this benchmark in order to determine whether the introduction of shared energy assets or scheduling strategies produces a meaningful improvement. In this way, the analysis does not simply describe isolated values from figures and tables, but explains why the observed results emerge and what they imply for the performance of the building energy system.

The discussion is developed mainly at building level, since the shared-energy

scenarios are designed to represent a collective residential system. However, selected archetype-level results are also reported where necessary to explain differences among household profiles and to clarify how family composition affects the distribution of energy use and cost. This distinction is particularly important because the aggregated building results reflect both the benefits of shared operation and the heterogeneity of the underlying apartment-level demand.

To improve readability, the chapter begins with a scenario matrix that summarizes the technological and operational features activated in each case. The results are then discussed in the same logical order as the scenario ladder itself: first the baseline benchmark, then the transition from Scenario 0 to Scenario 1, followed by the additional contribution of scheduling in Scenario 2, and finally the thermal comparison developed in Scenario 3. A short summary at the end of the chapter synthesizes the main findings and prepares the ground for the final conclusions of the thesis.

6.2 PV power forecasting

An hourly photovoltaic (PV) power series for the reference year 2025 was generated using an XGBoost regression model trained on PVGIS hourly PV power targets and meteorological predictors obtained from Open-Meteo. The forecasting workflow is reported only briefly here because its detailed implementation has already been defined in Chapter 5 and is documented in full in Appendix A. In summary, the model uses historical PVGIS outputs as the target variable, combines them with hourly weather and irradiance predictors, and produces a complete hourly forecast for the simulation year 2025. This forecast is required because the optimization framework operates at hourly resolution over the full year and therefore needs a physically consistent PV availability profile rather than a single annual yield value.

The predictor set includes irradiance-related variables such as global tilted irradiance (GTI), global horizontal irradiance (GHI), diffuse horizontal irradiance (DHI), and direct normal irradiance (DNI), together with additional meteorological features such as ambient temperature, wind speed, cloud-cover indicators, and calendar-based encodings of seasonal and diurnal position. The dataset is split chronologically into training (2018–2021), validation (2022), and test (2023) subsets. Model performance is evaluated on the 2023 test year, including a daylight-hours subset, in order to verify that the learned mapping is able to reproduce the PVGIS reference output with sufficiently high fidelity before it is applied to the forecast year 2025.

In the test evaluation, the model achieves strong agreement with the PVGIS reference series during daylight hours, indicating that the dominant physical relationship between irradiance conditions and PV output is captured with good accuracy. The purpose of the present section, however, is not to restate performance metrics in isolation, but to interpret what the forecast diagnostics imply for the subsequent optimization scenarios. In this thesis, the forecasting step is important not as an end in itself, but because it determines the hourly renewable input that drives the

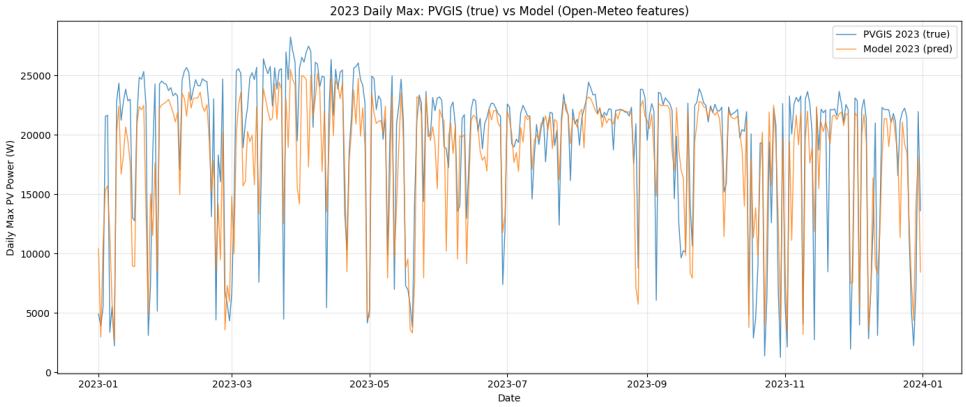

Figure 6.1: Daily peak PV power in 2023: PVGIS reference versus model prediction.

economic and operational interaction between rooftop PV generation, shared battery storage, and load scheduling in the following scenarios.

Figure 6.1 provides the first validation check by comparing the daily maximum PV power of the 2023 PVGIS reference series with the corresponding daily maximum predicted by the model on the 2023 test year. This figure is not intended to assess every hourly fluctuation in detail; rather, it is used to verify whether the model reproduces the annual envelope of peak PV production and its seasonal progression. The comparison shows that the predicted daily peaks follow the overall shape of the reference series closely across the year, with high peak values concentrated in the high-irradiance months and much lower values during winter. This is an important result because it indicates that the forecast does not suffer from an obvious seasonal distortion such as systematic overprediction in summer or underprediction in winter.

At the same time, the deviations that remain between the predicted and reference daily maxima are also meaningful. They are typically more visible on days with rapidly changing atmospheric conditions, when short-term cloud variability can produce sharp local peaks that are more difficult for a regression model to reproduce exactly. In other words, the model captures the dominant seasonal and meteorological structure of PV production well, while still smoothing part of the day-to-day irregularity associated with highly variable sky conditions. This is an acceptable and expected behaviour for the present application, since the optimization model mainly requires a realistic hourly PV availability signal with correct magnitude and seasonal timing, rather than perfect reconstruction of every extreme peak event.

A second and distinct diagnostic is provided by Figure 6.2, which compares the daily maximum PV power of the 2023 PVGIS reference year with the daily maximum values predicted for the forecast year 2025. The purpose of this figure is different from that of Figure 6.1. It is not a strict validation plot, because the 2025 series is a forecast input rather than an observed reference year. Instead, it is used to assess whether the forecasted 2025 PV input remains seasonally consistent and physically plausible for use in the optimization model.

The predicted 2025 daily maxima preserve the expected annual pattern of solar

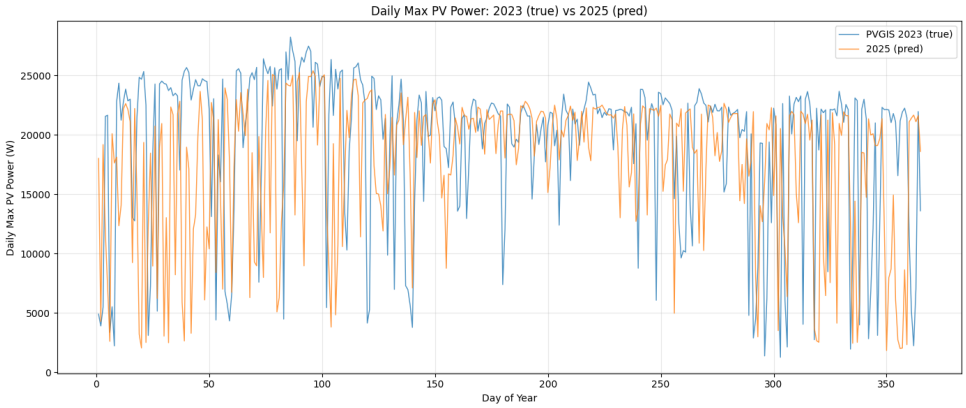

Figure 6.2: Daily peak PV power: 2023 PVGIS reference versus 2025 model prediction.

availability, with higher output concentrated in spring and summer and lower peaks during winter. This broad seasonal structure is the main feature that the optimization model needs in order to simulate the temporal interaction between local PV production, electricity demand, battery charging, and grid exchange. Visible differences between the 2023 PVGIS series and the 2025 predicted series should therefore not be interpreted as a contradiction of the forecasting approach. The two profiles correspond to different meteorological years, so some divergence is naturally expected as a result of inter-annual weather variability. What matters for the thesis is that the 2025 forecast remains realistic in magnitude, seasonal timing, and day-of-year structure, so that it can serve as a credible exogenous renewable input for the scenario analysis.

Taken together, Figures 6.1 and 6.2 support the use of the forecasted PV series in the optimization stage. The first figure shows that the model is able to reproduce the peak-envelope behaviour of the historical reference year with good seasonal consistency, while the second figure shows that the forecast year used in the simulation preserves a physically credible seasonal structure. The final forecast output is exported as an hourly time series for 2025 (column `P_pred` with timestamps) and is used to define the exogenous PV input profile P_t^{pv} in Scenarios 1–3. In this way, the following results are driven by a full-year hourly PV input that is not only technically available, but also checked for plausibility before being introduced into the shared-energy optimization framework.

6.3 Scenario setup and evaluation

Before discussing the numerical results, it is useful to summarize the technological and operational features activated in each scenario. Table 6.1 presents the scenario matrix adopted in this study. The matrix clarifies the progressive structure of the scenario ladder and shows how each scenario builds on the previous one. Scenario 0 represents the baseline case, with no PV, no BESS, no service scheduling, and no thermal intervention. Scenario 1 introduces shared PV and BESS. Scenario 2 extends this configuration by adding service scheduling. Scenario 3 represents the most complete configuration, in which the previous electricity-side measures are retained and the thermal dimension is also included.

This structure defines the evaluation logic of the chapter. Scenario 0 provides the benchmark against which the following scenarios are assessed. Scenario 1 is compared against Scenario 0 to quantify the effect of shared PV and BESS, Scenario 2 is compared against Scenario 1 to isolate the additional contribution of service scheduling, and Scenario 3 is interpreted as the full configuration. In this way, the results are discussed as a progressive transition rather than as disconnected simulations.

Table 6.1: Scenario structure

Scenario	PV	BESS	Service scheduling	Thermal
S_0	×	×	×	×
S_1	✓	✓	×	×
S_2	✓	✓	✓	×
S_3	✓	✓	✓	✓

6.4 Scenario 0: Baseline benchmark

Question addressed by Scenario 0

Scenario 0 establishes the benchmark operating condition of the case-study building before any technological or operational intervention is introduced. The purpose of this scenario is to answer the following baseline question:

> What is the annual energy procurement cost and grid-dependence profile of the building when no shared energy asset, scheduling strategy, or thermal retrofit measure is introduced?

Answering this question is essential because all subsequent scenarios are evaluated relative to the same service-preserving reference condition. Scenario 0 therefore provides the benchmark annual procurement cost, the benchmark pattern of grid dependence, and the benchmark gas-cost level from which the economic and operational effects of the later scenarios are measured.

Table 6.2: Scenario 0 building KPIs.

Indicator	Value	Unit
Annual electricity consumption	87,885.67	kWh
Electricity procurement cost	10,588.71	€
Annual gas consumption	12,562.11	Smc
Annual gas energy	134,414.62	kWh
Gas procurement cost	5,239.73	€
Total annual cost C_0	15,828.44	€

Table 6.3: Scenario 0 KPIs by archetype.

Archetype	Electricity (kWh)	El. cost (€)	Gas (kWh)	Gas cost (€)	Total cost (€)
Working couple	3,880.49	469.61	4,606.00	178.97	648.58
Family with 1 child	5,715.19	685.55	7,363.30	286.89	972.44
Family with 3 children	5,879.00	709.42	10,491.70	408.98	1,118.40
Retired couple	2,102.46	253.16	4,422.00	173.11	426.27

Scenario 0 represents the conventional operating condition of the residential building. Electricity demand is supplied entirely through grid imports under the hourly time-of-use tariff, while gas demand is met through the existing gas-based thermal configuration. No shared photovoltaic generation is installed, no shared battery energy storage system is available, no shiftable-service scheduling is activated, and no thermal retrofit measure is introduced. Electricity and gas demands are therefore treated as fully exogenous inputs, and the reported annual costs arise solely from the interaction between fixed service demand and the adopted electricity and gas price signals.

At building level, the baseline annual electricity consumption is 87,885.67 kWh, associated with an electricity procurement cost of EUR 10,588.71. Annual gas consumption reaches 12,562.11 Smc, equivalent to 134,414.62 kWh_{fuel}, with a gas procurement cost of EUR 5,239.73. The total annual baseline procurement cost is therefore EUR 15,828.44. These values define the reference point used in the following sections to quantify the cost savings and operational effects of the subsequent scenarios.

To support a more detailed interpretation, Table 6.3 reports the same baseline results by household archetype. This breakdown is important because the benchmark is clearly heterogeneous across dwelling types, and later savings cannot be interpreted as if all households started from the same demand and cost structure.

Image 6.3 summarizes the same benchmark graphically. The ranking is clear: the *Family with 3 children* archetype has the highest annual total cost, followed by *Family with 1 child*, *Working couple*, and *Retired couple*. This pattern is consistent with the

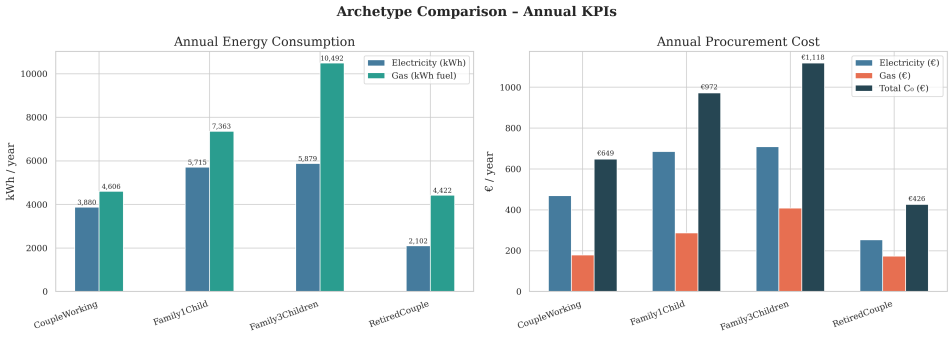

Figure 6.3: Scenario 0 KPIs by archetype.

Figure 6.4: Scenario 0 electricity by time-of-use band.

expected service demand of larger households. Higher occupancy generally increases appliance use, domestic hot-water demand, and overall thermal-service requirements, so both electricity and gas costs rise accordingly. The benchmark therefore already shows that household composition is a relevant driver of annual procurement cost.

A second key result concerns exposure to the time-of-use electricity tariff. Image 6.4 shows the distribution of annual electricity imports and electricity procurement cost across the F1, F2, and F3 bands. The baseline building imports 35,829 kWh in F1, 28,027 kWh in F2, and 24,029 kWh in F3. The corresponding annual costs are € 4,653.0, € 3,376.3, and € 2,559.4, respectively. The important point is not only that F1 is the most expensive band, but that the cost share is more concentrated in F1 than the energy share. This means that the economic burden of electricity procurement depends not only on how much electricity is purchased, but also on when it is purchased. This tariff structure is central to the later scenarios: PV, storage, and scheduling become valuable because they can reduce imports specifically during the most expensive hours.

The operational structure of the electricity baseline is clarified further by the building-level load-duration image shown in Figure 6.5. The sorted hourly profile shows that the highest electricity imports occur during a relatively small number of hours, whereas most of the year is characterized by substantially lower demand levels.

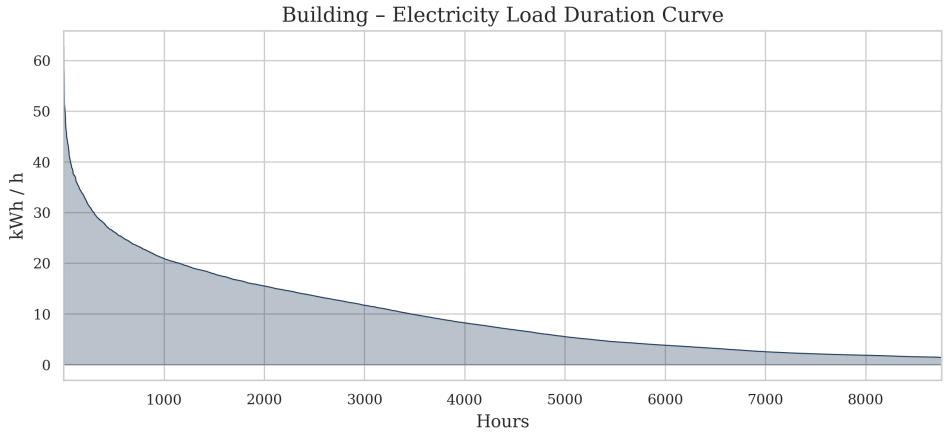

Figure 6.5: Scenario 0 electricity load-duration curve.

In other words, the baseline is peaky rather than flat. This is typical of aggregated residential demand: diversity across households smooths part of the annual profile, but simultaneous end-use activity still creates short periods of high demand. These peak hours matter disproportionately because they can coincide with expensive tariff periods, and they are exactly the hours in which shared PV, shared BESS, and shiftable-service scheduling can deliver the greatest benefit.

The thermal side of the benchmark can be interpreted more clearly by examining the gas demand structure in the same load-duration format. Image 6.6 shows that the sorted gas-demand profile declines much more sharply than the corresponding electricity profile and approaches zero for a substantial portion of the year. This indicates that gas demand is concentrated in a relatively limited number of hours dominated by space-heating and domestic-hot-water requirements, rather than being sustained uniformly across the annual horizon. In other words, the thermal side of the building is not only large in annual magnitude, but also structurally seasonal. This distinction is important for interpreting the scenario ladder: Scenarios 1 and 2 primarily improve the timing and sourcing of electricity demand, whereas the thermal problem addressed in Scenario 3 is driven by a different demand pattern, characterized by concentrated heating-season loads and long low-demand periods outside the main winter months.

Figure 6.7 complements this interpretation by showing the seasonal breakdown of electricity and gas consumption. Electricity demand is comparatively flatter across seasons than gas demand, reflecting the persistent contribution of residential appliance, lighting, and occupancy-driven uses throughout the year. Gas demand, by contrast, is strongly seasonal, with markedly higher values in winter and lower values in summer. This confirms that the baseline building is not purely an electricity problem. Even though Scenarios 1 and 2 primarily target electricity procurement through shared PV, storage, and scheduling, they do so on top of a thermal-demand background that remains structurally important in the annual energy balance.

Overall, Scenario 0 establishes a transparent benchmark for the rest of the chapter. It quantifies the original annual procurement cost of the case-study building, identifies

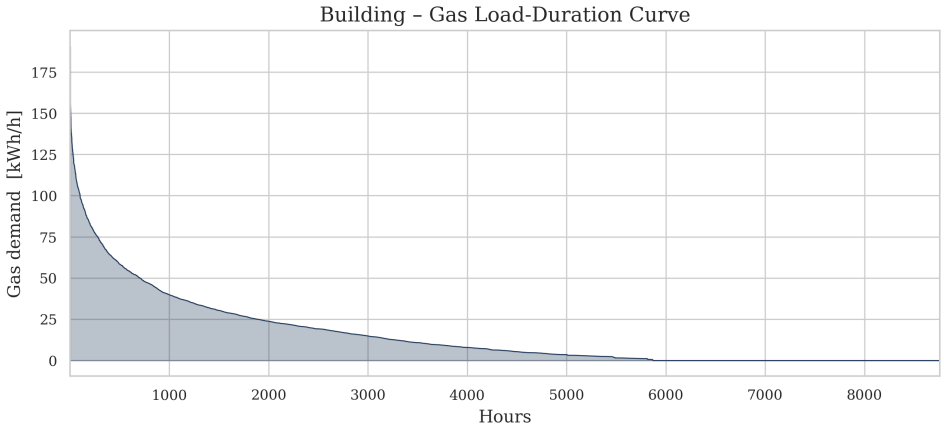

Figure 6.6: Scenario 0 gas load-duration curve.

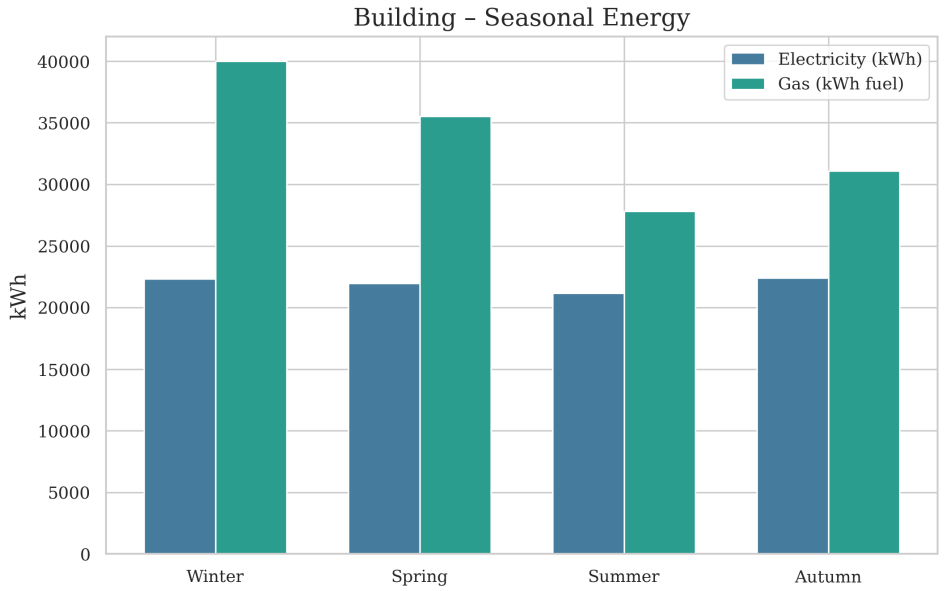

Figure 6.7: Scenario 0 seasonal energy use.

clear heterogeneity across household archetypes, shows that tariff exposure makes the timing of electricity imports economically important, and demonstrates that the electricity and gas sides of the building have different structural demand patterns. The baseline is therefore not only costly, but also operationally uneven: Electricity demand is year-round and peaky, whereas gas demand is strongly seasonal and concentrated in the heating period. These characteristics explain why the following scenarios are evaluated incrementally rather than in isolation.

The seasonal interpretation of the baseline can be refined further by decomposing monthly electricity and gas demand across the household archetypes. While the previous figure shows the aggregate building-level seasonal pattern, the following two figures clarify how that pattern is distributed across the representative household types.

Figure 6.8 shows that electricity demand remains present throughout the year and exhibits only moderate seasonal variation compared with the gas case discussed next.

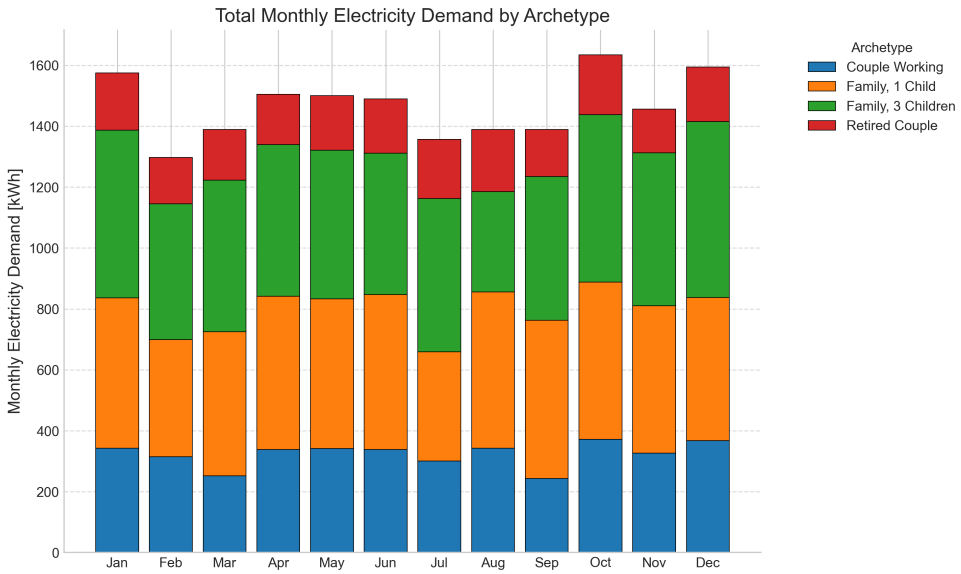

Figure 6.8: Monthly electricity demand by archetype in Scenario 0.

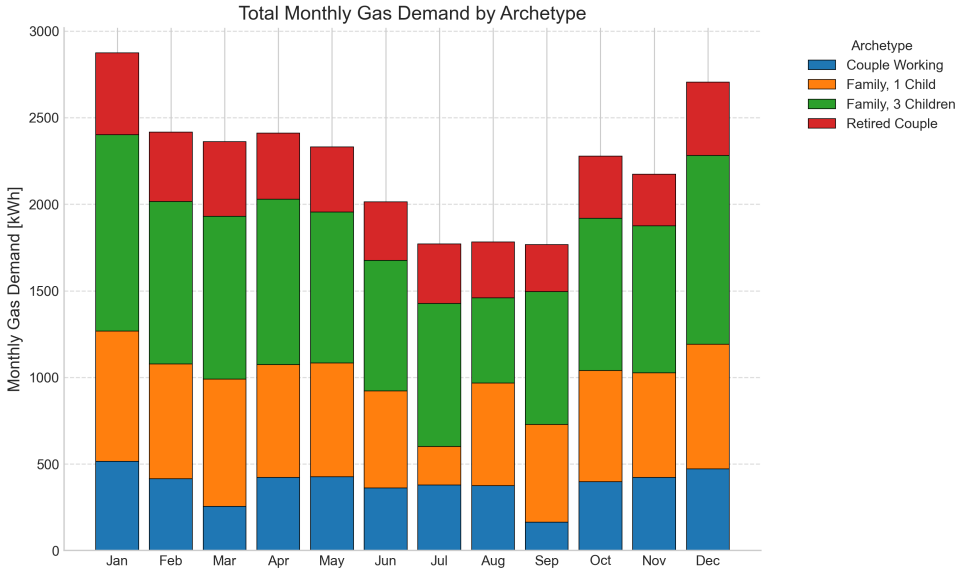

Figure 6.9: Monthly gas demand by archetype in Scenario 0.

This indicates that the electricity side of the building is structurally persistent rather than concentrated in a narrow seasonal window. The decomposition also shows that the two family archetypes contribute the largest share of monthly electricity demand in almost every month, confirming that larger households drive a disproportionate part of the building-level electricity requirement. The main analytical point is therefore not only that electricity demand is uneven across archetypes, but also that this heterogeneity remains stable across the year. From the perspective of the later scenarios, this is important because Scenarios 1 and 2 act on an electricity demand profile that remains relevant in all months, not only during a short seasonal period.

Figure 6.9 reveals a much stronger seasonal signature. Total monthly gas demand is clearly highest in the colder months and falls substantially during summer, which

confirms that the thermal side of the building is governed primarily by space-heating and domestic-hot-water requirements rather than by uniform year-round use. The family archetypes again account for the largest share of the total demand, but the more important result is the shape of the profile itself: unlike electricity, gas demand is strongly concentrated in the heating season. This distinction matters because it shows that the thermal problem later addressed in Scenario 3 is not simply a scaled version of the electricity problem. Instead, it has a structurally different temporal pattern, with pronounced winter concentration and much lower relevance in the warmer months.

Taken together, Figures 6.8 and 6.9 strengthen the interpretation of Scenario 0 by showing that the two energy carriers differ not only in annual magnitude, but also in their monthly structure across household types. Electricity remains comparatively persistent over the year, whereas gas exhibits a much stronger seasonal concentration. This confirms that the baseline building combines a year-round electricity procurement problem with a distinctly seasonal thermal-demand problem.

Mini-conclusion

Scenario 0 provides the service-preserving benchmark of the case-study building under conventional operation. In the absence of shared PV, shared BESS, shiftable-service scheduling, and thermal retrofit, the building records an annual electricity procurement cost of EUR 10,588.71, a gas procurement cost of EUR 5,239.73, and a total annual baseline cost of EUR 15,828.44. The baseline also shows clear heterogeneity across household archetypes, significant exposure to the high-price F1 tariff band, a peaky electricity-demand structure, and a strongly seasonal gas-demand profile. These characteristics define the reference state against which the economic and operational value of the subsequent scenarios is assessed.

6.5 Scenario 1: Shared PV and shared battery operation

Question addressed by Scenario 1

Scenario 1 addresses the following question: *after establishing the baseline procurement cost and grid dependence in Scenario 0, how much can annual electricity procurement cost and annual grid import be reduced by introducing a shared rooftop photovoltaic (PV) system and a shared battery energy storage system (BESS) at building level?* This scenario also separates the effect of *PV alone* from the additional effect of *shared storage*, so that the contribution of each intervention can be interpreted transparently. Since no service scheduling is introduced at this stage, any reduction in cost or imports is attributable only to shared PV generation and shared battery dispatch under the same time-of-use tariff structure used in the baseline. This is therefore a service-preserving comparison.

Reference to the baseline case

Scenario 0 is the benchmark for evaluating Scenario 1. In the baseline case, all building electricity demand is supplied by grid imports, no local PV generation is available, no battery is installed, and delivered services remain unchanged. At building level, the baseline annual electricity procurement cost is € 10,588.7/yr, the baseline annual grid import is 87,885.7 kWh/yr, and the total annual energy procurement cost including gas is € 15,828.4/yr. These values define the reference against which the economic and physical contribution of shared PV and shared storage must be measured.

Electricity-cost reduction: PV-only versus PV+BESS

Figure 6.10 compares annual electricity procurement cost for the three relevant cases: the baseline reference (Scenario 0), the PV-only subcase, and the full PV+BESS case. The baseline annual electricity procurement cost is € 10,588.7/yr. When only shared PV is introduced, the annual electricity procurement cost falls to € 7,228.9/yr, corresponding to a saving of € 3,359.8/yr relative to the baseline. When the shared BESS is added, the annual electricity procurement cost decreases further to € 6,573.7/yr. The total electricity-cost saving therefore reaches € 4,015.0/yr relative to Scenario 0, while the incremental value of shared storage beyond PV-only is € 655.2/yr.

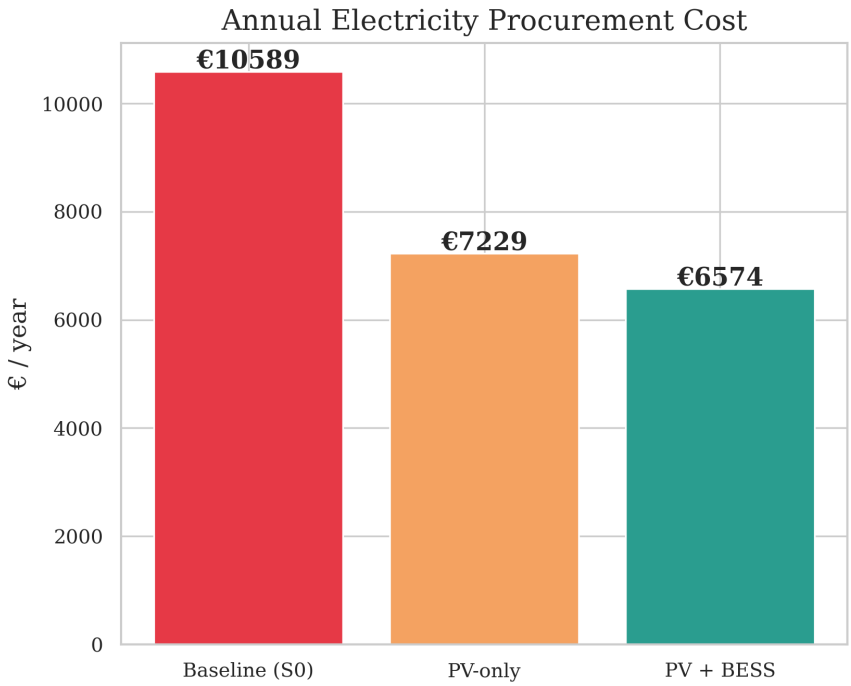

Figure 6.10: Annual total energy procurement cost in Scenario 1.

The same pattern is observed for the total annual energy procurement cost shown in Figure 6.11. Since gas demand and gas cost remain unchanged from Scenario 0, the reduction in total procurement cost is driven entirely by the electricity-side improvement. Total annual procurement cost decreases from € 15,828.4/yr in the baseline to € 12,468.7/yr in the PV-only case and to € 11,813.4/yr in the PV+BESS case. This confirms that the Scenario 1 benefit comes from improved electrical self-supply rather than from any change in gas consumption or thermal service demand.

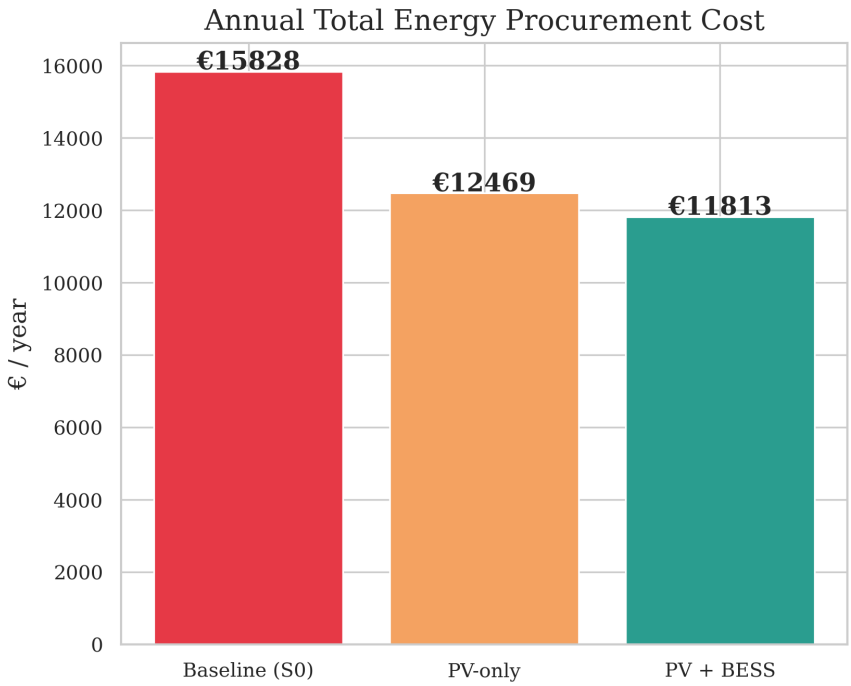

Figure 6.11: Annual total energy procurement cost in Scenario 1.

These results show a clear hierarchy of value. Shared PV provides the dominant first-step reduction because on-site generation directly offsets a substantial share of daytime electricity purchases. By contrast, the battery does not generate energy; its role is to improve the timing of energy use by storing part of the available surplus and releasing it later when the building still requires electricity. The contribution of storage is therefore real but incremental, which is exactly the pattern expected in a no-export behind-the-meter configuration.

Grid-import reduction and physical interpretation

Figure 6.12 reports annual grid import for the same three cases. In the baseline, annual grid import is 87,885.7 kWh/yr. With PV-only, annual grid import decreases to 61,017.8 kWh/yr, corresponding to a reduction of 26,867.9 kWh/yr. Adding shared BESS further reduces annual grid import to 56,334.2 kWh/yr, for a total reduction of 31,551.5 kWh/yr relative to the baseline. The incremental reduction enabled by shared storage beyond PV-only is therefore 4,683.6 kWh/yr.

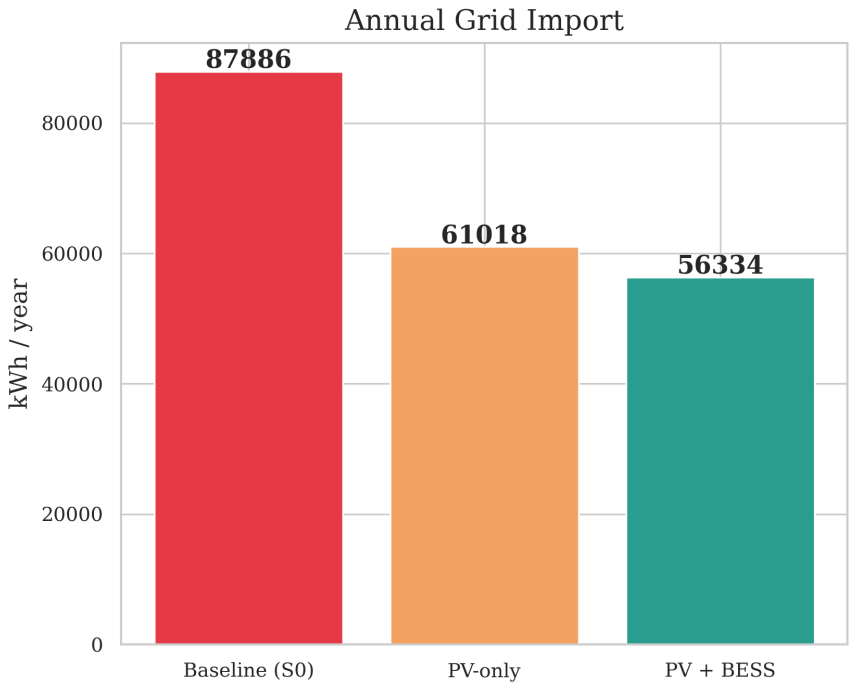

Figure 6.12: Annual grid import for the baseline, PV-only, and PV+BESS.

This result is important because it shows that the benefit of Scenario 1 is not only monetary but also physical. The lower electricity cost is accompanied by a substantial reduction in imported electricity, which means that the model is not simply shifting purchases between tariff hours; it is improving the actual matching between local generation, storage operation, and demand. At the same time, the result also shows that storage refines rather than dominates the PV effect. Most of the reduction is already achieved by PV alone, while the battery captures part of the remaining mismatch between midday PV production and later building demand. This is fully consistent with the Scenario 1 formulation, which reports both electricity cost and annual grid import as the core KPIs.

Battery operating signature

The battery operating signature provides further evidence of how the EMS uses storage. Annual discharged energy is approximately 8,403 kWh, corresponding to about 420 equivalent full cycles for a 20 kWh battery. This indicates that the battery is actively used across the year rather than remaining idle for long periods.

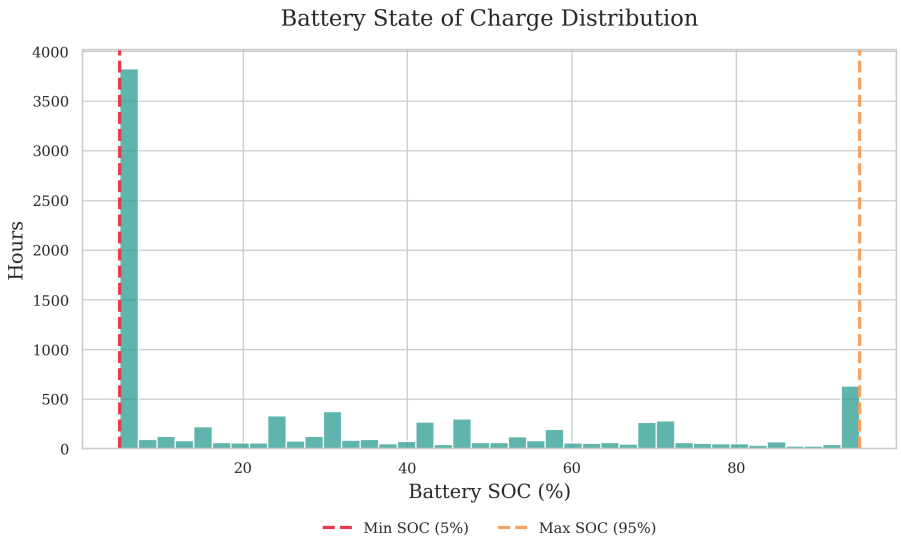

Figure 6.13: Distribution of hourly battery state of charge in Scenario 1.

Figure 6.13 shows that the battery spends a large share of the year close to its minimum operating bound of 5%, while higher SOC levels occur much less frequently. This pattern indicates that the shared battery is actively dispatched and often depleted soon after charging opportunities arise. In combination with the high annual cycling level, the SOC distribution suggests that the battery capacity is relatively limited compared with the magnitude and persistence of the aggregated building demand. Therefore, its role is not to eliminate grid dependence, but to improve the temporal matching between PV availability, imported electricity, and building load. :contentReference[oaicite:14]index=14

Mini-conclusion

Scenario 1 demonstrates that shared rooftop PV is the primary driver of cost reduction at building level, while shared battery storage provides a secondary but meaningful incremental benefit. Relative to the baseline, the PV-only case reduces annual electricity procurement cost from € 10,588.7 to € 7,228.9 and annual grid import from 87,885.7 to 61,017.8 kWh. Adding shared BESS further reduces electricity procurement cost to € 6,573.7 and annual grid import to 56,334.2 kWh. The battery creates this additional value by increasing the utilization of on-site PV generation, reducing curtailment, and improving the temporal matching between energy availability and building demand. Since no service scheduling is introduced in this scenario, all reported gains are achieved in a service-preserving manner and therefore represent the pure contribution of shared PV and shared storage dispatch under the adopted no-export framework.

Table 6.4: Annual electricity cost and grid import from the baseline to Scenario 2.

Configuration	Annual electricity cost (€)	Annual grid import (kWh)
Baseline	10,589	87,886
PV only	7,229	61,018
PV + BESS	6,574	56,334
PV + BESS + shiftable loads	6,040	53,475

6.6 Scenario 2: Shared PV + BESS + Shiftable

Scenario question and analytical focus

Scenario 2 addresses the following question: *to what extent can demand-side scheduling of selected shiftable electrical services further improve the performance of the shared PV+BESS configuration introduced in Scenario 1?*

The purpose of this scenario is not merely to show that flexibility reduces annual electricity cost, but to explain *how* this reduction is achieved. In particular, Scenario 2 examines whether the scheduling of shiftable household services can:

1. increase the direct coincidence between building demand and PV generation,
2. reduce residual grid import after the introduction of the battery,
3. improve the utilization of locally generated PV electricity, and
4. provide additional economic benefit beyond the PV+BESS case.

The optimization preserves the same annual service demand as in the baseline schedule and therefore does not create artificial savings by suppressing appliance use. Instead, the model repositions eligible loads within their admissible operating windows so that consumption is better aligned with PV-rich hours and with the battery dispatch strategy. In this way, the scenario isolates the operational value of *load shifting* on top of the energy-system assets already introduced in Scenario 1.

Annual scenario comparison

Table 6.4 compares the annual building-level electricity cost and annual grid import across the four electricity configurations considered in Chapters 3 and 4: the baseline case without shared assets, the PV-only case, the PV+BESS case, and the full PV+BESS+shiftable-load case.

The progression across the four configurations is structurally coherent. The introduction of shared PV already delivers the largest single improvement, because a significant fraction of daytime electricity demand can be covered directly by on-site generation. Relative to the baseline, the PV-only case reduces annual electricity cost by about 3,360 €/year. The addition of the battery produces a second, smaller improvement, since excess PV generation can be shifted from production hours to

later hours with positive residual demand. Finally, the introduction of shiftable loads yields a further reduction of about 534 €/year relative to the PV+BESS case, bringing the total annual electricity cost down to 6,040 €. This corresponds to an overall reduction of about 43 % relative to the baseline configuration.

A similar pattern appears in annual grid import. Shared PV reduces grid purchases substantially because locally generated electricity directly offsets daytime demand. The battery then decreases imports further by storing part of the PV surplus and discharging it when PV production is unavailable. Scenario 2 adds another reduction of about 2,859 kWh/year relative to the PV+BESS case, and the total decrease with respect to the baseline reaches about 34,411 kWh/year, or approximately 39 %. The key point is that flexibility does not replace the role of PV or storage; instead, it increases their operational effectiveness by moving part of the controllable demand into hours where locally generated or previously stored energy is available.

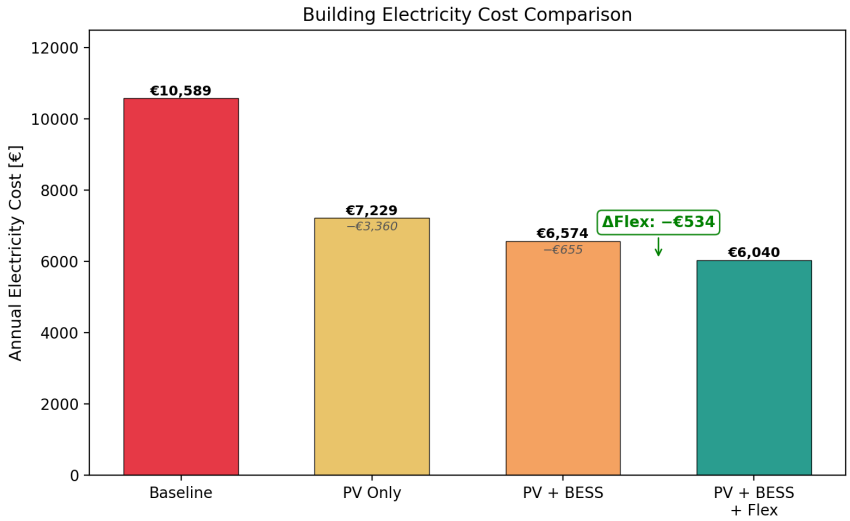

Figure 6.14: Annual electricity cost across the four electricity configurations.

Figure 6.14 shows that the economic contribution of flexibility is incremental rather than dominant. This is an important result. It means that the value of load shifting is real, but bounded by the size of the shiftable demand, by comfort and scheduling constraints, and by the fact that PV and storage have already captured the largest part of the easy savings. Therefore, Scenario 2 should be interpreted as a refinement of the PV+BESS strategy rather than as a fundamentally different cost driver.

Figure 6.15 confirms the same hierarchy from an energy-balance perspective. The difference between the PV+BESS and PV+BESS+shiftable-load cases is smaller than the difference between the baseline and the PV cases, but it is still meaningful because it shows that flexibility improves local matching between demand and available energy resources. In other words, Scenario 2 does not reduce annual demand; it reduces the *timing mismatch* between demand and supply.

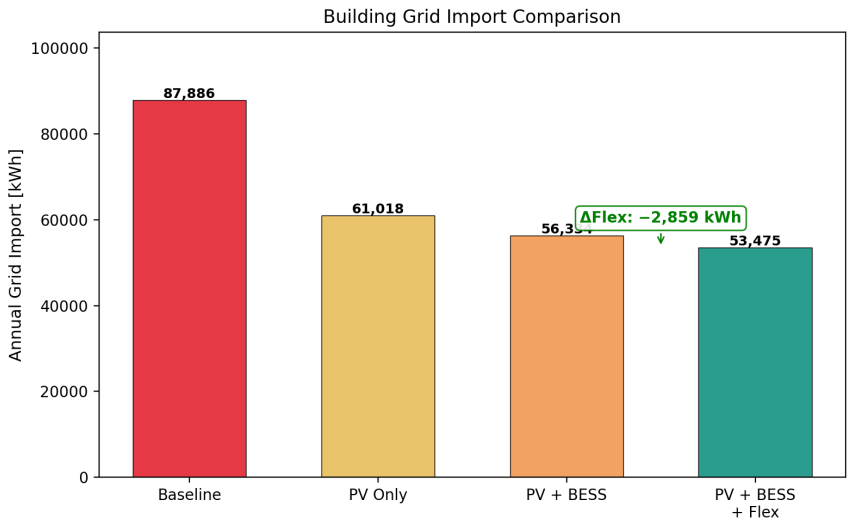

Figure 6.15: Annual grid import across the four electricity configurations.

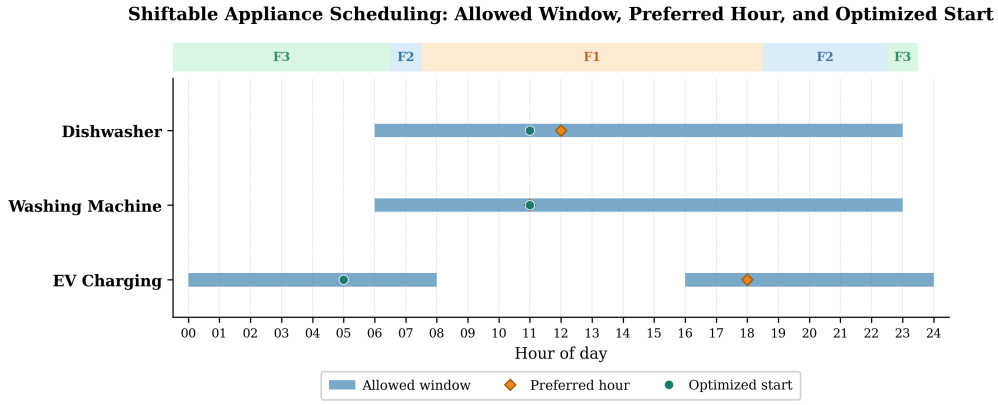

Figure 6.16: Allowed window, preferred hour, and optimized start.

How the load is shifted

The effect of load shifting in Scenario 2 is better understood by examining the scheduling logic of the flexible services rather than only comparing aggregate start-time distributions. Figure 6.16 therefore reports, for representative shiftable appliances, the admissible operating window, the preferred hour, and the optimized start selected by the model. This figure is intended as a representative illustration of the scheduling mechanism rather than as a complete behavioural description of every household archetype. In practice, preferred operating times differ across household types, especially between working and non-working occupants, but the same optimization principle applies in all cases: each service can only be shifted within its feasible window and only when the resulting relocation improves the overall operating condition of the system.

Figure 6.16 shows that the optimizer does not move all loads uniformly, nor does it shift them exclusively toward the hours of maximum PV generation. Instead, each appliance is repositioned selectively within its admissible window according to

the combined influence of PV availability, tariff conditions, battery interaction, and residual grid dependence. This point is important because it clarifies that Scenario 2 is not based on arbitrary schedule manipulation, but on feasible and system-oriented rescheduling.

The washing machine remains essentially aligned with its preferred hour, indicating that its original timing is already close to an acceptable operating point and that the additional benefit of moving it further would be limited. The dishwasher undergoes only a modest adjustment, shifting slightly earlier while still remaining well inside its feasible window. This suggests that the optimizer applies only a targeted correction when the expected system benefit is relatively small. By contrast, EV charging experiences the strongest relocation, moving from an evening-preferred start toward the early-morning window. This is the clearest example of flexibility being exploited where it provides the greatest operational value. The fact that this optimized start is not colocated with the midday PV peak confirms that the scheduling decision is not driven by PV alone. Rather, the optimizer seeks the most favourable feasible hour from the perspective of the overall system, including the avoidance of more critical import periods later in the day.

The figure should also be interpreted with some behavioural caution. A single preferred hour cannot represent all household archetypes equally well, since working households, families, and retired occupants typically exhibit different daily routines. For example, appliance use in working households is more likely to be concentrated in the evening than in midday hours. Nevertheless, this does not invalidate the figure. Its role is not to reproduce every behavioural variant in detail, but to show the scheduling principle adopted in Scenario 2: user preference is respected where possible, flexibility is used only within admissible bounds, and stronger shifts occur only for services whose wider temporal freedom makes them more valuable for system coordination.

Overall, Figure 6.16 explains why Scenario 2 improves upon the PV+BESS configuration without relying on unrealistic demand relocation. The benefit arises from selective and feasible schedule adjustments, with the largest changes assigned to the services that can move most effectively and with the least behavioural penalty. The system-level impact of this redistribution is examined next through the comparison of the average 24-hour grid-import profiles across Scenarios 0, 1, and 2.

The system-level consequence of this rescheduling is shown in Figure 6.17, which compares the average hourly grid-import profiles of Scenario 0, Scenario 1, and Scenario 2, together with the average PV generation profile and the hourly change in net grid import caused by load shifting.

Figure 6.17 shows that the benefit of load shifting does not come from reducing grid import uniformly across all hours. Instead, Scenario 2 redistributes grid dependence over the day. Relative to Scenario 1, the optimized schedule increases net import during the early-morning hours, particularly in the off-peak period, and reduces it from late morning through the evening. This pattern is consistent with the scheduling logic observed previously: highly flexible demand, especially EV charging,

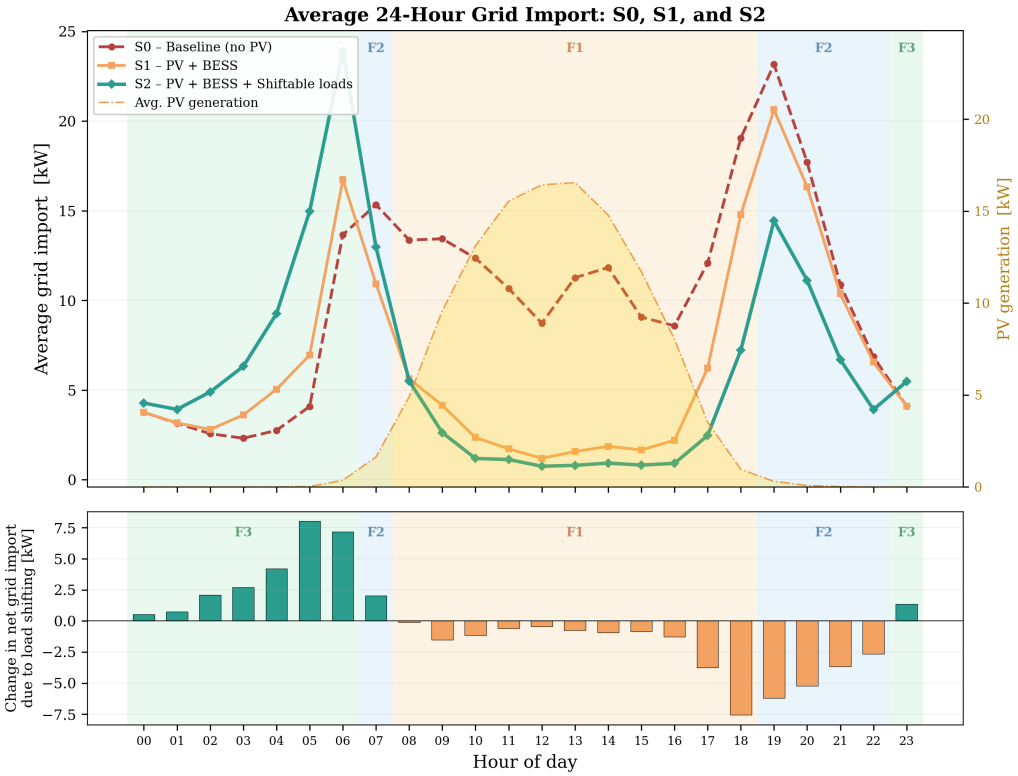

Figure 6.17: Average 24-hour grid import and PV generation in Scenarios 0–2.

is moved away from the evening preference and into earlier hours where the system can accommodate it more effectively. The most important result is the reduction in residual grid import during the late afternoon and evening, where the PV+BESS case still exhibits a pronounced dependence on the grid. This means that the operational value of Scenario 2 lies not in suppressing demand, but in repositioning flexible consumption toward hours that improve the overall coordination between PV availability, battery operation, and grid exchange. In other words, the optimizer uses temporal flexibility to relieve the most critical import periods rather than simply shifting demand at random.

Taken together, Figure 6.16 and 6.17 explain why Scenario 2 improves upon the PV+BESS case. The first image shows that the schedule changes are feasible and selective, while the second shows that these changes translate into a more favorable daily grid- import profile. The representative summer and winter weeks discussed next provide a more detailed temporal view of how this coordinated behaviour unfolds under different seasonal operating conditions.

Representative summer and winter dispatch

Representative weekly dispatch plots help interpret how the optimized schedule interacts with the shared PV system and battery across seasons. Figures 6.18 and 6.19 show the optimized building load, PV generation plotted as negative power, residual grid import, battery discharge, and battery state of charge over a summer week and a winter week, respectively.

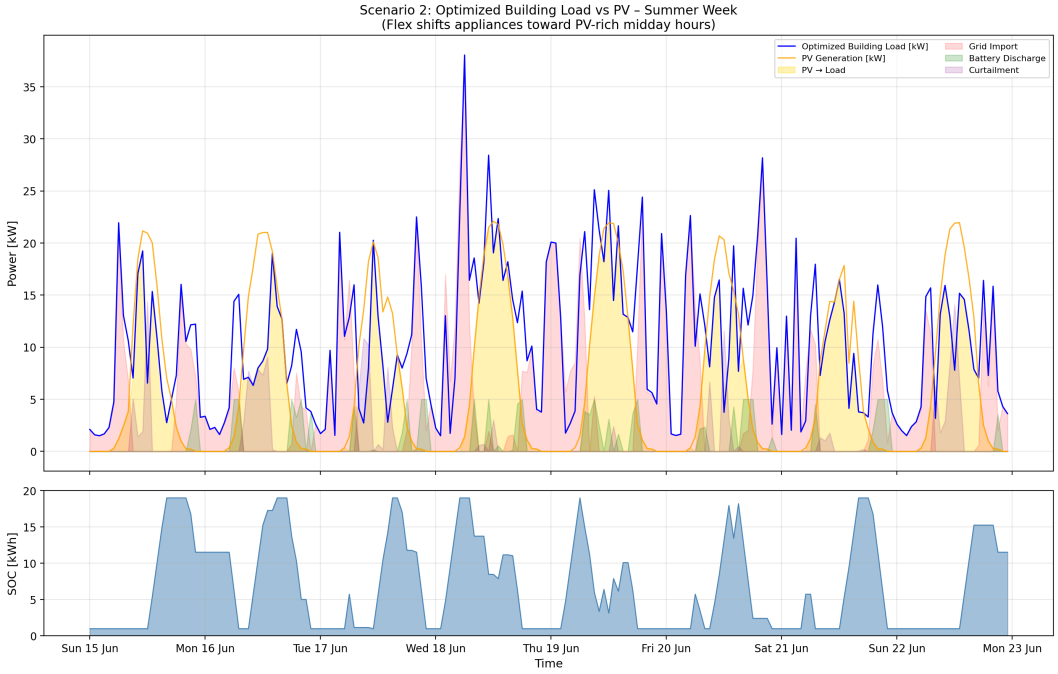

Figure 6.18: Representative summer-week dispatch in Scenario 2.

In summer, the effect of load shifting is more visible because PV production is high and sustained during midday hours. The optimized load curve tends to rise in periods where solar generation is abundant, which indicates that part of the shiftable demand has been intentionally moved to coincide with local production. As a result, daytime PV surplus is more effectively absorbed by the building, reducing both curtailment and the need for later grid purchases. The battery then complements this behavior by storing part of the residual surplus and discharging during subsequent hours of positive demand. In practical terms, the summer result shows that flexibility increases the *direct-use value* of PV before the battery even intervenes.

In winter, the same mechanism is present but less pronounced. PV production is lower and more intermittent, so the optimization has fewer high-generation hours to exploit. Consequently, the role of load shifting becomes more constrained: some demand can still be moved toward the limited daytime PV window, but the potential for direct solar self-consumption is smaller than in summer. Under these conditions, the battery remains useful, yet it also has less surplus PV energy available for charging. This seasonal contrast explains why annual flexibility gains are moderate rather than dramatic. The scheduling logic is effective, but its benefit is fundamentally dependent on the magnitude and timing of PV availability.

For the battery state of charge, the technical operating range should be interpreted relative to the nominal 20 kWh capacity. The minimum and maximum allowed SOC are 1 kWh and 19 kWh, corresponding to 5 % and 95 %, respectively. In the final thesis version, the SOC axis should therefore be shown in percentage rather than in kWh, since this communicates the actual operating window more clearly and directly.

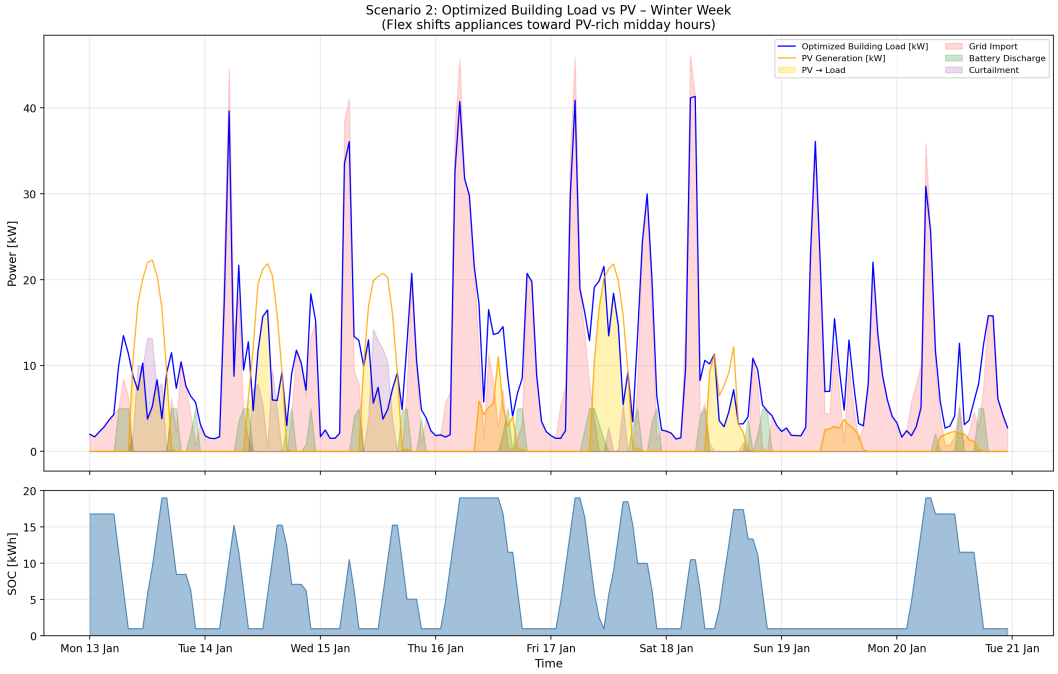

Figure 6.19: Representative winter-week dispatch in Scenario 2.

Interpretation of the incremental shiftable benefit

The annual results demonstrate that the incremental contribution of flexibility is positive but moderate. This is the correct result, not a weakness of the model. The baseline-to-PV transition captures the large structural gain of replacing imported daytime electricity with local generation. The PV-to-PV+BESS transition captures the next layer of value by transferring part of the solar surplus over time. Scenario 2 then operates on the remaining mismatch, which is narrower and more constrained.

This explains why the additional savings from shiftable loads are smaller in magnitude than the savings produced by PV or by PV+BESS. The optimization is acting on a restricted portion of total demand, and that demand must still satisfy service timing, comfort, and technical feasibility. In other words, Scenario 2 improves the system mainly through *better temporal coordination*, not through large changes in annual energy quantity.

The Scenario 2 results therefore support three conclusions. First, load shifting is operationally valuable because it reduces residual grid dependence even after PV and storage are installed. Second, its benefit is strongest when PV availability is high enough to provide attractive hours for rescheduling, which makes the effect more visible in summer than in winter. Third, the economic value of flexibility is complementary to shared energy assets, not independent of them. The model shows that flexibility works best when it is integrated with PV and storage rather than analyzed as an isolated measure.

Scenario 2 mini-conclusion

Scenario 2 answered the question of whether the addition of shiftable electrical loads can improve the shared PV+BESS configuration at building level. The results show that it can. Relative to the PV+BESS case, load shifting reduces annual electricity cost by about 534 € and annual grid import by about 2,859 kWh. Relative to the baseline, the full PV+BESS+shiftable-load configuration reduces electricity cost by about 43 % and grid import by about 39 %.

More importantly, the scenario clarifies the mechanism behind these gains. The optimization does not lower service demand; it reschedules eligible loads toward hours with stronger PV availability and more favorable system conditions. This increases direct PV self-consumption, improves the interaction between demand and storage, and decreases the residual need for imported electricity. The benefit remains moderate because only part of the load is shiftable and because realistic operational constraints are preserved. Nevertheless, the results confirm that demand scheduling is a meaningful complementary strategy once shared PV and battery storage are already in place.

6.7 Scenario 3: Thermal retrofit optimization

Question addressed by Scenario 3

Scenario 3 evaluates the thermal retrofit dimension of the case study by comparing two alternative technologies for supplying the same hourly thermal-service demand: a condensing gas boiler and an air-source heat pump (ASHP). Unlike Scenarios 0–2, which focus on electricity procurement and operational coordination, the purpose of Scenario 3 is to answer a narrower but important techno-economic question:

> Which thermal supply technology can provide the required space-heating and domestic-hot-water services at lower annualized cost under the same thermal demand conditions?

This distinction is important for the logic of the chapter. In the previous scenarios, the main economic question concerns electricity procurement under different shared energy configurations. In Scenario 3, the comparison is instead investment-aware and service-preserving. The thermal demand itself is not reduced, and the result depends on the trade-off between annual operating cost and annualized capital burden. The scenario therefore isolates the effect of changing the thermal conversion technology without mixing it with the shared-PV, shared-BESS, or service-scheduling effects analyzed earlier.

Archetype-level annual thermal cost comparison

The primary economic result is the comparison of annualized thermal-supply cost across the four household archetypes. Figure 6.20 compares the existing gas-boiler case with the ASHP retrofit for the same annual thermal-service requirement.

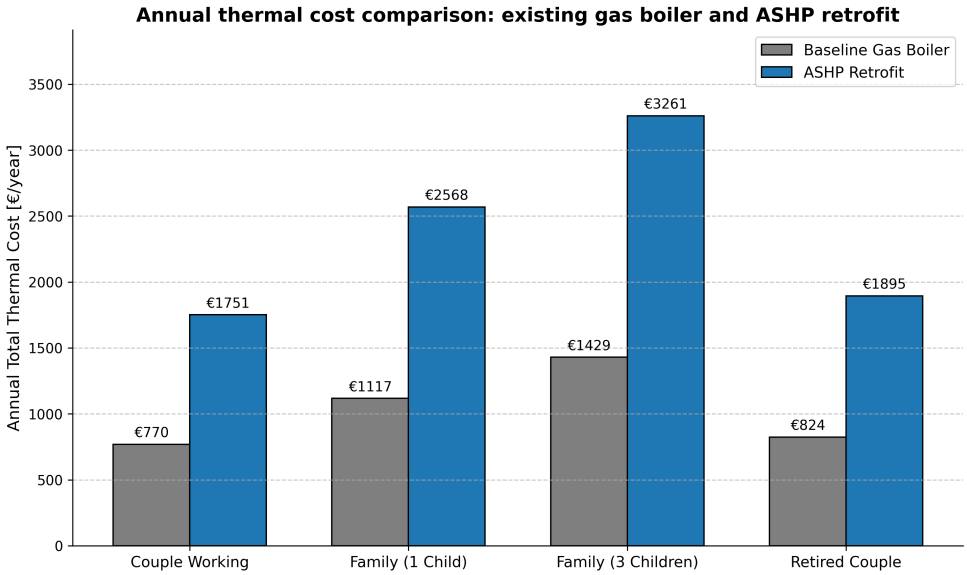

Figure 6.20: Annual thermal cost comparison: existing gas boiler and ASHP retrofit.

Table 6.5: Scenario 3 final optimization KPIs.

	Thermal demand [kWh$_{th}$/yr]			Peak load	Gas boiler [€/yr]		ASHP [€/yr]	
Archetype	Q_{SH}	Q_{DHW}	Q_{tot}	[kW$_{th}$]	OPEX	Total	OPEX	Total
Couple working	4300.43	3915.11	8215.54	9.71	353.21	769.50	299.96	1750.87
Family (1 child)	5375.54	6258.81	11634.34	14.37	501.22	1117.27	421.28	2568.44
Family (3 children)	6181.87	8917.91	15099.78	18.15	651.11	1429.42	547.88	3260.59
Retired couple	4837.98	3758.66	8596.64	10.58	370.69	824.32	313.64	1894.72

Figure 6.20 shows that, under the assumptions adopted in this thesis, the gas-boiler option remains cheaper than the ASHP retrofit for all archetypes. The result is not marginal for the three non-retired archetypes, where the cost gap is large enough to indicate a clear preference for the boiler under the present parameterization. The economic conclusion is therefore straightforward: within the current CAPEX, O&M, lifetime, discount-rate, and fixed-COP assumptions, the ASHP does not recover its higher annualized investment cost through operating-cost savings alone.

To support the interpretation of this comparison, Table 6.5 reports the main thermal-service and cost indicators by archetype, including annual space-heating demand, annual domestic-hot-water demand, total useful thermal demand, required design capacity, and the annualized total-cost outputs of the two technology options.

The table confirms that annualized cost increases broadly with useful thermal demand, but that demand magnitude alone does not explain the technology choice. The key result is that the ASHP remains more expensive even though it is thermodynamically more efficient. This indicates that the final decision is being driven by the structure of costs rather than by service demand alone.

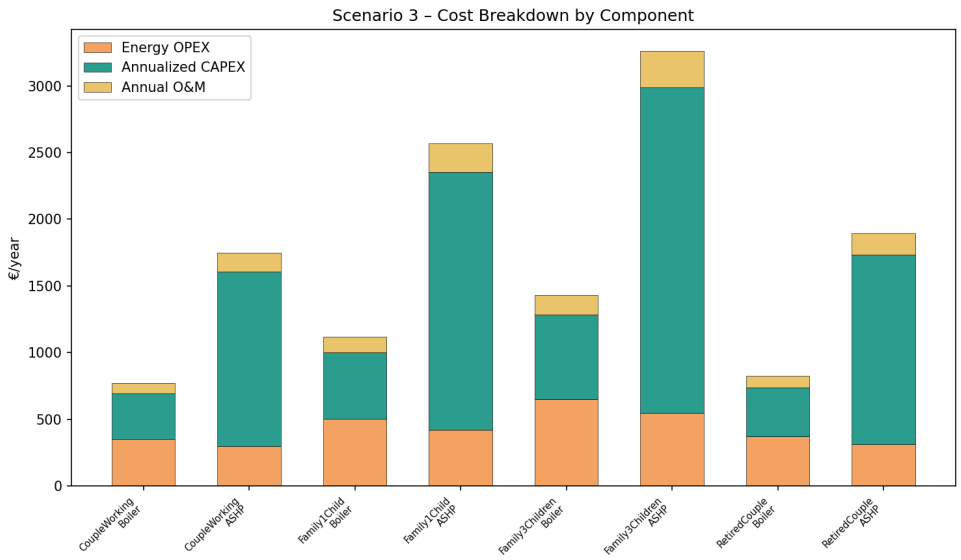

Figure 6.21: Stacked annual cost breakdown by archetype and technology.

Cost-component decomposition

The mechanism behind the solver's decision is clarified by Figure 6.21, which decomposes annual cost into energy OPEX, annualized CAPEX, and annual O&M for each archetype and technology.

Figure 6.21 explains why the boiler is selected even though the ASHP is more efficient in delivering useful heat. The ASHP consistently benefits from lower energy-related operating cost, because its electricity input is lower than the boiler's fuel input after thermal conversion losses are taken into account. In purely operational terms, this is the expected and physically consistent result. However, the ASHP also carries a much larger annualized CAPEX burden than the gas boiler for all archetypes, and its annual O&M term is higher as well because it is modeled as a fraction of a larger upfront investment. Under the present assumptions, this capital-cost penalty dominates the operating-cost benefit, which is why the total annualized cost of the ASHP remains higher.

Retrofit benefit indicator

The investment relevance of this result is made explicit by Figure 6.22, which reports the annual cost penalty of the ASHP relative to the existing gas boiler.

Figure 6.22 translates the cost-breakdown result into direct decision language. Positive values indicate that the ASHP retrofit increases annualized cost relative to the baseline boiler case, whereas negative values would indicate a net saving. In the present results, all four archetypes show a positive annual penalty, which means that the retrofit is not economically justified under the modeled assumptions. This is the most direct expression of the Scenario 3 conclusion, because it answers the practical investment question immediately: the ASHP is not only more expensive in total, but also fails to generate an annualized cost saving relative to the existing thermal

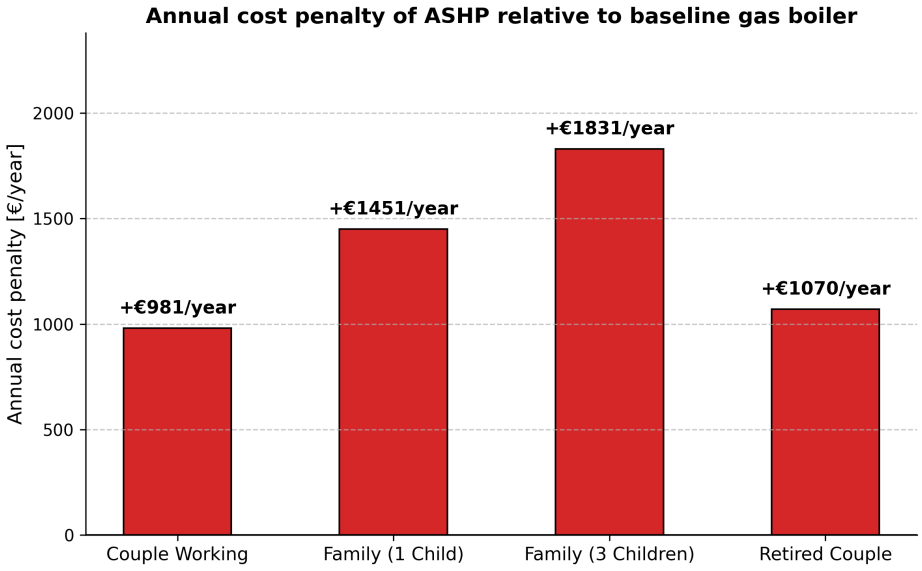

Figure 6.22: Annual ASHP cost penalty relative to the baseline gas boiler.

supply option.

The largest penalty is observed for the *Family (3 children)* archetype, followed by *Family (1 child)*, *Retired couple*, and *Couple working*. This ranking is consistent with the larger design capacities and thermal-service requirements associated with the higher-demand households. At the same time, the uniformly positive bars confirm that the result is not an isolated edge case. Under the current assumptions, the economic disadvantage of the ASHP persists across all archetypes.

Interpretation of the retrofit outcome

Taken together, Figures 6.20, 6.21, and 6.22 provide a coherent interpretation of the Scenario 3 outcome. The first figure identifies which technology is cheaper, the second explains the internal economic mechanism behind that result, and the third converts the same outcome into an explicit investment signal. This combination is important because it prevents an ambiguous reading of the cost breakdown. The ASHP is not rejected because it performs poorly as a heat-supply technology; it is rejected because, under the present economic assumptions, its efficiency advantage is not large enough to compensate for its higher annualized capital burden.

This conclusion should therefore be interpreted as assumption-dependent rather than universal. Scenario 3 does not show that ASHP retrofits are inherently uneconomic in all contexts. Instead, it shows that under the specific CAPEX values, component lifetimes, O&M treatment, discount-rate assumptions, and fixed-COP representation adopted in this thesis, the annual operating-cost savings are insufficient to offset the higher annualized investment cost of electrified thermal supply.

Mini-conclusion

Scenario 3 answered the question of whether replacing the existing gas boiler with an ASHP reduces the annualized cost of supplying the same space-heating and domestic-hot-water services. Under the assumptions adopted in this thesis, it does not. For all four household archetypes, the gas-boiler option remains the lower-cost choice once annualized CAPEX, annual O&M, and energy-related operating cost are combined in a single yearly objective function. The retrofit-benefit indicator makes this result directly visible: all archetypes exhibit a positive annual cost penalty for the ASHP relative to the baseline boiler case. The base-case recommendation of Scenario 3 is therefore to retain the gas boiler under the modeled assumptions, while recognizing that this conclusion is sensitive to investment cost, lifetime, discount rate, and heat-pump performance assumptions.

Chapter 7

Conclusion and Future Prospects

7.1 Conclusion

This thesis developed a service-preserving optimization framework to evaluate how shared photovoltaic generation, shared battery storage, flexible demand scheduling, and thermal retrofit options affect the economic and operational performance of a multi-family residential building. The central contribution is not a single technology model in isolation, but a structured decision-support methodology that combines hourly data preparation, annual techno-economic optimization, and staged scenario comparison within one coherent framework.

The work addressed a practical residential decision problem: how a multi-unit building can be modernized while preserving the same underlying household services and while keeping the economic interpretation of each intervention transparent. To answer this question, the thesis used a scenario-ladder structure in which each scenario isolates a distinct decision layer. This separation is important because it allows the value of local generation, storage, scheduling, and thermal technology choice to be interpreted progressively rather than being mixed into one opaque final result.

The first main outcome of the thesis is the establishment of a clear and traceable baseline benchmark. Scenario 0 showed that, under conventional operation and in the absence of shared PV, shared BESS, flexible scheduling, and thermal technology substitution, the case-study building records an annual electricity procurement cost of EUR 10,588.71, a gas procurement cost of EUR 5,239.73, and a total annual baseline cost of EUR 15,828.44. The benchmark analysis also showed that the building is heterogeneous across household archetypes, that exposure to the high-price F1 tariff band is economically important, and that electricity and gas follow different structural patterns, with electricity remaining comparatively persistent across the year and gas being much more seasonal. These baseline results are important because they define the reference state against which all later interventions are interpreted.

The second main outcome concerns the value of shared rooftop PV and shared battery storage at building level. Scenario 1 showed that shared PV is the dominant source of electricity-cost reduction, while shared BESS provides a secondary but still meaningful incremental benefit. Relative to the baseline, the PV-only case reduces

annual electricity procurement cost from EUR 10,588.7 to EUR 7,228.9 and annual grid import from 87,885.7 kWh to 61,017.8 kWh. When shared BESS is added, the annual electricity procurement cost falls further to EUR 6,573.7 and annual grid import to 56,334.2 kWh. These results show that the battery does not create value independently of PV, but rather improves the temporal use of locally generated electricity by reducing mismatch between daytime PV production and later demand.

The third main outcome concerns the additional value of flexibility. Scenario 2 showed that flexible scheduling of selected electrical services improves performance beyond the PV+BESS configuration, but that this improvement is moderate rather than dramatic. The full PV+BESS+shiftable-load case reduces annual electricity procurement cost to EUR 6,040 and annual grid import to 53,475 kWh. This result is important not because it produces the largest absolute saving, but because it clarifies the role of flexibility within the hierarchy of interventions. Load shifting does not create value by reducing service demand. Instead, it improves temporal coordination by moving eligible consumption toward more favourable hours within feasible operating windows. Its contribution is therefore positive and meaningful, but naturally smaller than the structural gains already captured by shared PV and shared storage.

The fourth main outcome concerns the thermal retrofit comparison. Scenario 3 evaluated whether replacing the existing gas boiler with an air-source heat pump can reduce the annualized cost of supplying the same thermal-service demand. Under the assumptions adopted in this thesis, it cannot. For all household archetypes, the gas boiler remains the lower-cost option once annualized CAPEX, annual O&M, and energy-related operating cost are combined into one yearly objective. The cost breakdown showed that the ASHP benefits from lower operating cost, but that this advantage is insufficient to compensate for its higher annualized investment burden. The retrofit-benefit indicator made this conclusion explicit by showing a positive annual ASHP cost penalty for all archetypes. The Scenario 3 result is therefore clear: under the present assumptions, the ASHP is physically credible but not economically preferred.

Taken together, these findings lead to a coherent overall conclusion. For the studied building, the most robust economic improvement is created by shared rooftop PV, strengthened by shared battery storage and complemented by realistic load shifting. By contrast, thermal electrification through ASHP does not become cost-optimal under the present annualized-cost assumptions. This means that the modernization value of the building is not concentrated in one universal intervention. Instead, it emerges from a hierarchy of decision layers, each of which contributes differently and must be interpreted on its own terms.

From a methodological perspective, the thesis also contributes a transparent way of structuring residential-energy decisions. By separating the scenario ladder into distinct operational and investment layers, the framework avoids the ambiguity of one-step co-optimization in which the source of value can become difficult to interpret. The thesis therefore contributes not only numerical results for one case-study building,

but also a decision-oriented modeling logic that makes it possible to identify where value is created, why it is created, and under which assumptions it remains valid.

Overall, the thesis answered its research question by showing that a service-preserving, scenario-based optimization framework can provide transparent annual techno-economic decision support for residential building modernization. For the present case study, the strongest and most robust pathway is based on shared PV, reinforced by shared battery storage and improved through realistic scheduling of flexible demand, while the ASHP retrofit remains economically unattractive under the current assumptions. The framework developed in this work should therefore be understood not only as a simulation tool, but also as a structured method for interpreting residential modernization choices step by step.

7.2 Future Prospects

The framework developed in this thesis provides a strong basis for expansion toward larger and more integrated residential energy systems. The most important future direction is not simply the refinement of the present building model, but the extension of the same optimization logic from a single multi-family building to a broader residential energy-community scale. This is the natural next step, because many of the benefits associated with distributed generation, shared storage, and coordinated flexibility become even more relevant when multiple buildings are allowed to interact within a common local energy system.

A first and particularly promising extension is therefore the transition from the single-building case study to a multi-building or community-scale framework. In such a setting, the optimization could coordinate shared PV generation, centralized or distributed battery storage, flexible household demand, and electric-vehicle charging across several buildings rather than within one building boundary only. This would make it possible to evaluate how diversity across buildings, user routines, and demand profiles can further improve local self-consumption, reduce peak grid dependence, and strengthen the collective economic value of shared energy assets.

A second important prospect concerns the inclusion of broader community-level coordination mechanisms. Future developments could incorporate peer-to-peer energy sharing, collective battery allocation rules, community EV-charging coordination, and alternative electricity-sharing arrangements between residential users. Such an extension would shift the framework from a building optimization tool toward a more comprehensive local-energy-management methodology suitable for residential energy communities.

A third promising direction is the refinement of behavioural representation. The present thesis already models flexibility in a service-preserving way, which provides a strong foundation for more detailed behavioral treatment in future work. Archetype-specific preferences, occupancy schedules, and differentiated flexibility windows could be represented more explicitly so that working households, families with children, and retired occupants are captured with greater realism. This would further strengthen

the practical relevance of scheduling results, especially in larger community-scale applications.

Another relevant extension concerns the market and policy environment in which the framework operates. Future studies could examine the same optimization structure under dynamic electricity tariffs, export remuneration, collective self-consumption rules, investment subsidies, tax incentives, or carbon-related fuel penalties. This would make it possible to assess how the economic ranking of PV, storage, flexibility, and thermal retrofit options changes under different regulatory conditions and would increase the usefulness of the framework for planning and policy-oriented analysis.

The thermal dimension also offers substantial room for extension. In particular, the Scenario 3 methodology can be expanded beyond the present gas-boiler-versus-ASHP comparison toward a richer community-level thermal strategy that includes temperature-dependent heat-pump performance, hybrid systems, thermal storage, and interaction between electrified heating and shared PV generation. This would allow future work to evaluate not only annualized retrofit economics, but also the operational synergies between local renewable electricity and electrified thermal supply.

A further prospect is the inclusion of environmental criteria alongside the economic assessment. Future applications of the framework could combine annualized cost, self-consumption, grid-import reduction, and carbon-emission performance within a multi-criteria decision structure. This would be especially valuable at energy-community scale, where economic and environmental objectives are often pursued simultaneously rather than separately.

Finally, the methodology can be replicated across a wider range of case studies, including different climates, tariff regimes, building typologies, and residential user compositions. Such extensions would make it possible to assess which conclusions are robust across contexts and which remain specific to the present case-study building. In this sense, the thesis should be seen not only as the analysis of one residential building, but also as the foundation of a scalable methodology for future research on shared residential energy systems and local energy communities.

In summary, the most important future perspective of this work is its expansion from a single-building optimization framework to a coordinated multi-building residential energy-community framework. By scaling the methodology in this direction, the same service-preserving and decision-oriented logic developed in this thesis can be applied to a broader and more impactful class of residential energy systems.

Appendix A

Photovoltaic system sizing and PV time-series generation

This appendix reports the PV sizing inputs and the detailed PV time-series generation workflow used to construct the hourly PV series PV_t for 2025.

A.1 PV sizing inputs

The PV array is constrained by rooftop geometry. The usable roof area for PV deployment is $A_{\text{roof}} = 642$ m^2 and the installed layout corresponds to $N_{\text{mod}} = 127$ PV modules. The installed peak capacity is computed as:

$$P_{\text{PV}}^{\text{inst}} = N_{\text{mod}} \cdot P_{\text{mod}}, \tag{A.1}$$

where P_{mod} is the module nameplate power from the selected module datasheet. The geometric feasibility condition is:

$$N_{\text{mod}} \cdot A_{\text{mod}} \leq A_{\text{roof}}, \tag{A.2}$$

where A_{mod} is the module area. Tilt/azimuth and loss assumptions are kept consistent between PVGIS configuration and Open-Meteo GTI query, as required by the adopted PVGIS-consistent workflow [51, 53].

0.49

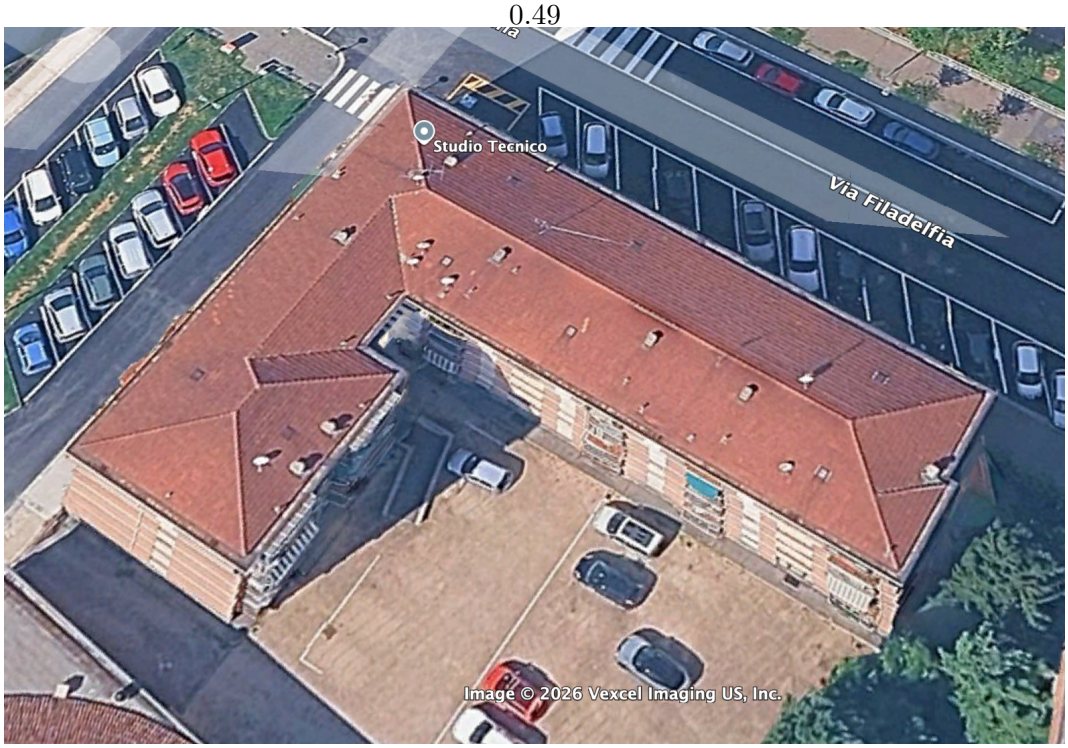

Figure A.1: Original rooftop view.

0.49

Figure A.2: PV layout used for sizing.

Figure A.3: Rooftop geometry and PV module layout adopted in the case study.

A.2 PV time-series generation workflow

The PV time-series generation for 2025 follows these steps:

1. **PVGIS target series (2018–2023).** Retrieve PVGIS hourly PV power

output (PPP) for 2018–2023 under a fixed PV system configuration (capacity, tilt/azimuth, losses) [51].

2. **Open-Meteo predictors (2018–2025).** Retrieve hourly meteorological and radiation predictors for 2018–2025 from Open-Meteo Historical Weather API, including GHI/DHI/DNI and GTI for the same coordinates [52, 53].

3. **Timestamp alignment.** Since Open-Meteo radiation variables are past-hour means, shift Open-Meteo timestamps by 30 minutes to represent midpoint values and merge to PVGIS timestamps using nearest-neighbor matching within a small tolerance [53].

4. **Train/validation/test split.** Split the paired dataset chronologically: train (2018–2021), validation (2022), test (2023).

5. **Model training.** Train an XGBoost regressor to predict PVGIS PPP from Open-Meteo predictors [55].

6. **Evaluation.** Compute MAE/RMSE on 2023 (and optionally daylight-only).

7. **Inference for 2025.** Apply the trained model to Open-Meteo 2025 predictors to produce PV_t for all 8,760 hours.

Model hyperparameters and additional diagnostic plots (optional) can be reported here if required by the thesis format.

Appendix B

Supplementary data preparation details

This appendix collects supplementary details that support reproducibility but would interrupt the flow of Chapter 3.

B.1 national holidays for TOU assignment

National holidays are treated as off-peak (F3) for all 24 hours in the tariff mapping used in this thesis. Easter Sunday and Easter Monday are computed algorithmically for 2025.

B.2 LoadProfileGenerator configuration summary

LoadProfileGenerator (LPG) is used to generate full-year hourly residential electricity demand profiles for the four household archetypes [38]. Table B.2 summarizes the key simulation settings used for all archetypes.

Table B.1: Italian national holidays used for time-band assignment (2025).

Date (2025)	Holiday
Jan 1	New Year's Day
Jan 6	Epiphany
Apr 20	Easter Sunday
Apr 21	Easter Monday
Apr 25	Liberation Day
May 1	Labour Day
Jun 2	Republic Day
Aug 15	Assumption (Ferragosto)
Nov 1	All Saints' Day
Dec 8	Immaculate Conception
Dec 25	Christmas Day
Dec 26	St. Stephen's Day

Table B.2: LPG simulation configuration

Setting	Value
Simulation period	01/01/2025–31/12/2025
Time step	3600 s (1 hour)
Energy carrier	Electricity
Context	European residential routines/holidays
Output exported	total household electricity demand (kWh per hour)

Appendix C

Component lifetimes and replacement assumptions

This appendix documents the lifetime assumptions used to annualize capital expenditures (CAPEX) and, when relevant, to model replacement costs consistently with a one-year (2025) operational optimization horizon.

C.1 Economic horizon and annualization

All operational decisions are optimized over year 2025 with hourly resolution. Investment costs are converted to an equivalent annual cost to allow comparison on a /year basis.

For a component i with upfront CAPEX I_i (), lifetime L_i (years), and real discount rate r, the capital recovery factor (CRF) is:

$$\mathrm{CRF}(r, L_i) = \frac{r(1+r)^{L_i}}{(1+r)^{L_i} - 1}. \tag{C.1}$$

The annualized CAPEX of component i is:

$$C_{i,\mathrm{ann}}^{\mathrm{cap}} = I_i \cdot \mathrm{CRF}(r, L_i). \tag{C.2}$$

Component lifetimes and replacement assumptions

Lifetime values

Table C.1 reports the lifetimes adopted for energy-system assets used directly in Scenario 3. Table C.2 reports indicative household appliance lifetimes; these are only relevant if an appliance replacement decision is explicitly modeled and are otherwise reported for transparency.

Note: In the main formulation, these lifetimes are primarily used for CAPEX annualization. Explicit replacement decisions are only relevant when replacement is directly modeled or when the analysis horizon is sufficiently long.

Table C.1: Asset lifetimes used for CAPEX annualization in Scenario 3.

Asset	Lifetime L_i (y)	Repl. modeled?	Modeling use
PV modules (rooftop plant)	25	Optional	Annualized over PV lifetime.
PV inverter	10–15	Optional	Replacement only for long horizons.
BESS (battery pack)	10–15	Optional	Annualized using L_{bat}.
Air-source heat pump (ASHP)	15–20	Optional	CAPEX included if ASHP is selected.
Gas boiler (condensing)	13–21	Optional	CAPEX included if boiler is selected.

Table C.2: Indicative appliance lifetimes.

Appliance / component	Lifetime (y)	Modeling use
Dishwasher	9	Binary replacement decision if appliance replacement is included.
Dryer	13	Binary replacement decision if appliance replacement is included.
Washing machine	5–15	Use a base-case value, such as 10 years.
Water heater (electric or gas)	10–11	Relevant only if DHW equipment is modeled separately.

Note: Appliance lifetime values are indicative and should be interpreted as modeling assumptions rather than exact service-life guarantees.

Bibliography

[1] Tom Mike Terlouw, Tarek AlSkaif, Christian Bauer, Marco Mazzotti, and Russell McKenna. "Designing residential energy systems considering prospective costs and life cycle GHG emissions". In: *Applied Energy* 331 (2023), p. 120362. DOI: 10.1016/j.apenergy.2022.120362 (cit. on pp. 1–10, 12, 14).

[2] Iain Staffell, Stefan Pfenninger, and Nathan Johnson. "A global model of hourly space heating and cooling demand at multiple spatial scales". In: *Nature Energy* 8 (2023), pp. 1328–1344. DOI: 10.1038/s41560-023-01341-5 (cit. on pp. 1–3, 5).

[3] Ishanki De Mel, Floris Bierkens, Xinyao Liu, Matthew Leach, Mona Chitnis, Lirong Liu, and Michael Short. "A decision-support framework for residential heating decarbonisation policymaking". In: *Energy* 268 (2023), p. 126651. DOI: 10.1016/j.energy.2023.126651 (cit. on pp. 1–3, 5, 6, 8–10, 12, 14).

[4] Nikolaos G. Paterakis, Ozan Erdinc, and João P. S. Catalão. "An overview of demand response: Key-elements and international experience". In: *Renewable and Sustainable Energy Reviews* 69 (2017), pp. 871–891. DOI: 10.1016/j.rser.2016.11.167 (cit. on pp. 1, 2, 5–9).

[5] Marc Beaudin and Hamidreza Zareipour. "Home energy management systems: A review of modelling and complexity". In: *Renewable and Sustainable Energy Reviews* 45 (2015), pp. 318–335. DOI: 10.1016/j.rser.2015.01.046 (cit. on pp. 1, 2, 5–9).

[6] Pierluigi Mancarella. "MES (multi-energy systems): An overview of concepts and evaluation models". In: *Energy* 65 (2014), pp. 1–17. DOI: 10.1016/j.energy.2014.01.020 (cit. on pp. 1, 4).

[7] Michael J. Fell. "Energy services: A conceptual review". In: *Energy Research & Social Science* 27 (2017), pp. 129–140. DOI: 10.1016/j.erss.2017.02.010 (cit. on pp. 2, 4–7).

[8] Gerald Kalt, Dominik Wiedenhofer, Christoph Görg, and Helmut Haberl. "Conceptualizing energy services: A review of energy and well-being along the Energy Service Cascade". In: *Energy Research & Social Science* 53 (2019), pp. 47–58. DOI: 10.1016/j.erss.2019.02.026 (cit. on pp. 2, 4–7).

[9] Tobias Beck, Hendrik Kondziella, Guillaume Huard, and Thomas Bruckner. "Optimal operation, configuration and sizing of generation and storage technologies for residential heat pump systems in the spotlight of self-consumption of photovoltaic electricity". In: *Applied Energy* 188 (2017), pp. 604–619. DOI: 10.1016/j.apenergy.2016.12.041 (cit. on pp. 2–5, 7–10, 14).

[10] Sandro Schopfer, Verena Tiefenbeck, and Thorsten Staake. "Economic assessment of photovoltaic battery systems based on household load profiles". In: *Applied Energy* 223 (2018), pp. 229–248. DOI: 10.1016/j.apenergy.2018.03.185 (cit. on pp. 2, 3, 5, 8, 12).

[11] Simon I. Sun, Barnaby D. Smith, Richard G. A. Wills, and Alexander F. Crossland. "Effects of time resolution on finances and self-consumption when modeling domestic PV-battery systems". In: *Energy Reports* 6 (2020), pp. 157–165. DOI: 10.1016/j.egyr.2020.03.020 (cit. on pp. 2, 5, 6, 8, 12).

[12] Cameron A. Belton and Peter D. Lunn. "Smart choices? An experimental study of smart meters and time-of-use tariffs in Ireland". In: *Energy Policy* 140 (2020), p. 111243. DOI: 10.1016/j.enpol.2020.111243 (cit. on p. 3).

[13] Rasmus Luthander, Joakim Widén, Daniel Nilsson, and Jenny Palm. "Photovoltaic self-consumption in buildings: A review". In: *Applied Energy* 142 (2015), pp. 80–94. DOI: 10.1016/j.apenergy.2014.12.028 (cit. on pp. 3, 6, 8).

[14] Jacopo Torriti. "Price-based demand side management: Assessing the impacts of time-of-use tariffs on residential electricity demand and peak shifting in Northern Italy". In: *Energy* 44.1 (2012), pp. 576–583. DOI: 10.1016/j.energy.2012.05.043 (cit. on p. 3).

[15] Paolo Lazzeroni, Gianmarco Lorenti, Aldo Canova, and Maurizio Repetto. "Economic and environmental perspectives of flexible demand in PV-based Italian energy communities with residential end-users". In: *Energy and Buildings* 342 (2025), p. 115853. DOI: 10.1016/j.enbuild.2025.115853 (cit. on p. 4).

[16] Filip Nastić, Nebojša Jurišević, Danijela Nikolić, and Davor Končalović. "Harnessing open data for hourly power generation forecasting in newly commissioned photovoltaic power plants". In: *Energy for Sustainable Development* 81 (2024), p. 101512. DOI: 10.1016/j.esd.2024.101512 (cit. on pp. 6, 12).

[17] Nathalie Frieß, Ulrich Pferschy, David Raese, and Joachim Schauer. "Assessing the potential of forecast-based optimization in renewable energy communities with flexible electricity, heat and mobility resources". In: *Applied Energy* 401 (2025), p. 126664. DOI: 10.1016/j.apenergy.2025.126664 (cit. on pp. 6, 10, 12).

[18] Albert Molderink, Vincent Bakker, Matthijs G. C. Bosman, Johann L. Hurink, and Gerard J. M. Smit. "Management and Control of Domestic Smart Grid Technology". In: *IEEE Transactions on Smart Grid* 1.2 (2010), pp. 109–119. DOI: 10.1109/TSG.2010.2055904 (cit. on p. 7).

[19] Zdenek Bradač, Vaclav Kaczmarczyk, and Petr Fiedler. "Optimal Scheduling of Domestic Appliances via MILP". In: *Energies* 8.1 (2015), pp. 217–232. DOI: 10.3390/en8010217 (cit. on pp. 7, 9).

[20] Kin Cheong Sou, James Weimer, Henrik Sandberg, and Karl Henrik Johansson. "Scheduling Smart Home Appliances Using Mixed Integer Linear Programming". In: *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*. 2011, pp. 5144–5149. DOI: 10.1109/CDC.2011.6161081 (cit. on pp. 7, 9).

[21] Abubakar Sani Hassan, Liana Cipcigan, and Nick Jenkins. "Optimal battery storage operation for PV systems with tariff incentives". In: *Applied Energy* 203 (2017), pp. 422–441. DOI: 10.1016/j.apenergy.2017.06.043 (cit. on p. 8).

[22] Yujie Wang, Serena Rinaldi, et al. "Multi-objective home appliance scheduling with implicit and interactive user satisfaction modelling". In: *Applied Energy* 267 (2020), p. 114690. DOI: 10.1016/j.apenergy.2020.114690 (cit. on p. 9).

[23] Francesco Nicoletti, Mario Antonio Cucumo, and Natale Arcuri. "Cost optimal sizing of photovoltaic-battery system and air–water heat pump in the Mediterranean area". In: *Energy Conversion and Management* 270 (2022), p. 116274. DOI: 10.1016/j.enconman.2022.116274 (cit. on pp. 9, 10).

[24] Lei Zhang, Guohui Feng, Kailiang Huang, Yang Bi, Shasha Chang, and Ainong Li. "Design and optimization for photovoltaic heat pump system integrating thermal energy storage and battery energy storage". In: *Energy and Buildings* 329 (2025), p. 115277. DOI: 10.1016/j.enbuild.2025.115277 (cit. on pp. 9, 10).

[25] International Energy Agency Photovoltaic Power Systems Programme. *Trends in Photovoltaic Applications*. 2023 (cit. on p. 17).

[26] Rasmus Luthander, Joakim Widén, Daniel Nilsson, and Jenny Palm. "Photovoltaic self-consumption in buildings: A review". In: *Applied Energy* 142 (2015), pp. 80–94 (cit. on pp. 17, 18, 26).

[27] Peter Lund, Juuso Lindgren, Joni Mikkola, and Jyri Salpakari. "Review of energy system flexibility measures to enable high levels of variable renewable electricity". In: *Renewable and Sustainable Energy Reviews* 45 (2015), pp. 785–807 (cit. on pp. 17, 18, 21, 22).

[28] David Parra, Gordon Walker, and Mark Gillott. "Solar PV and battery storage systems for self-consumption in residential buildings". In: *Renewable and Sustainable Energy Reviews* 78 (2017), pp. 1322–1335 (cit. on pp. 17, 18, 26).

[29] REN21. *Renewables 2023 Global Status Report*. 2023 (cit. on p. 18).

[30] Joakim Widén and Ewa Wäckelgård. "A review of methods for modeling household electricity load profiles". In: *Energy and Buildings* 47 (2012), pp. 202–212 (cit. on p. 20).

[31] Ian Richardson, Murray Thomson, David Infield, and Conor Clifford. “Domestic electricity use: A high-resolution energy demand model”. In: *Energy and Buildings* 42.10 (2010), pp. 1878–1887 (cit. on p. 20).

[32] Peter Palensky and Dietmar Dietrich. “Demand side management: Demand response, intelligent energy systems, and smart loads”. In: *IEEE Transactions on Industrial Informatics* 7.3 (2011), pp. 381–388 (cit. on pp. 20, 23, 26).

[33] Pierluigi Siano. “Demand response and smart grids—A survey”. In: *Renewable and Sustainable Energy Reviews* 30 (2014), pp. 461–478 (cit. on pp. 20, 23, 26).

[34] David Connolly et al. “Heat Roadmap Europe: Combining district heating with heat savings to decarbonise the EU energy system”. In: *Energy Policy* 65 (2014), pp. 475–489 (cit. on pp. 21, 22).

[35] Iain Staffell and Daniel Brett. “A review of domestic heat pumps”. In: *Energy & Environmental Science* 5.11 (2012), pp. 9291–9306 (cit. on pp. 22, 23).

[36] Michael James Fell. “Energy services: A conceptual review”. In: *Energy Research & Social Science* 27 (2017), pp. 129–140. DOI: 10.1016/j.erss.2017.02.010 (cit. on pp. 39, 40).

[37] Abubakar Sani Hassan, Liana Cipcigan, and Nick Jenkins. “Optimal battery storage operation for PV systems with tariff incentives”. In: *Applied Energy* 203 (2017), pp. 422–441. DOI: 10.1016/j.apenergy.2017.06.043 (cit. on p. 40).

[38] Noah Pflugradt, Peter Stenzel, Leander Kotzur, and Detlef Stolten. “LoadProfileGenerator: An Agent-Based Behavior Simulation for Generating Residential Load Profiles”. In: *Journal of Open Source Software* 7.71 (2022), p. 3574. DOI: 10.21105/joss.03574. URL: https://doi.org/10.21105/joss.03574 (cit. on pp. 56, 58, 106).

[39] Noah Pflugradt, Peter Stenzel, Leander Kotzur, and Detlef Stolten. “LoadProfileGenerator: An Agent-Based Behavior Simulation for Generating Residential Load Profiles”. In: *Journal of Open Source Software* 7.71 (2022), p. 3574. DOI: 10.21105/joss.03574 (cit. on pp. 56, 58).

[40] OpenStreetMap contributors. *Researcher Information (citation and community guidance)*. Accessed 2026-02-14. URL: https://wiki.openstreetmap.org/wiki/Researcher_Information (cit. on p. 58).

[41] *UNI/TS 11300-2:2014 – Energy performance of buildings – Part 2: Evaluation of primary energy need and system efficiencies for space heating and domestic hot water production*. Tech. rep. Italian technical specification referenced for DHW energy balance formulation. UNI – Ente Italiano di Normazione, 2014 (cit. on p. 62).

[42] Jay Burch and Craig Christensen. *Towards Development of an Algorithm for Mains Water Temperature*. Tech. rep. Available via OSTI: 981988. National Renewable Energy Laboratory (NREL), 2007 (cit. on p. 62).

[43] *Technical Support Document: Energy Efficiency Program for Consumer Products and Commercial and Industrial Equipment – Water Heaters*. Tech. rep. Provides representative Energy Factor values for gas tankless and condensing water heaters. National Renewable Energy Laboratory (NREL), 2013 (cit. on p. 62).

[44] ARERA. *Cosa è il potere calorifico superiore (P) indicato in bolletta?* ARERA – Atlante per il consumatore (gas billing guide). Explains PCS (higher heating value) usage on Italian gas bills. accessed 2026-03-04 (cit. on p. 62).

[45] N. G. Paterakis, O. Erdinc, I. N. Pappi, A. G. Bakirtzis, and J. P. S. Catalao. "Coordinated Operation of a Neighborhood of Smart Households Comprising Electric Vehicles, Energy Storage and Distributed Generation". In: *IEEE Transactions on Smart Grid* 7.6 (2016), pp. 2736–2747 (cit. on p. 63).

[46] S. Lee and D.-H. Choi. "Two-level distributed deep reinforcement learning for optimal energy management of a smart home". In: *Sensors* 20.2157 (2020). DOI: 10.3390/s20082157 (cit. on p. 63).

[47] ARERA. *Cosa sono i prezzi per fasce? (F1/F2/F3)*. https://www.arera.it/atlante-per-il-consumatore/elettricita/il-mercato-dellenergia/le-offerte-commerciali/cosa-sono-i-prezzi-per-fasce. Web page, accessed 2026-02-14. n.d. (Cit. on p. 64).

[48] ARERA. *Gas – La bolletta: Cosa si paga*. n.d. URL: https://www.arera.it/atlante-per-il-consumatore/gas/la-bolletta/cosa-si-paga (visited on 02/14/2026) (cit. on p. 65).

[49] ARERA. *Valore CMEM,m – Servizio di tutela della vulnerabilità*. n.d. URL: https://www.arera.it/en/comunicati-operatore/dettaglio/valore-cmemm-servizio-di-tutela-della-vulnerabilita (visited on 02/14/2026) (cit. on p. 65).

[50] ARERA. *Gas – Quota fissa: cosa è e quando si paga*. n.d. URL: https://www.arera.it/atlante-per-il-consumatore/gas/la-bolletta/cose-la-quota-fissa (visited on 02/14/2026) (cit. on p. 65).

[51] European Commission and Joint Research Centre. *PVGIS 5.3 release notes*. Web page. Accessed: 2026-02-13. URL: https://joint-research-centre.ec.europa.eu/photovoltaic-geographical-information-system-pvgis/pvgis-releases/pvgis-53_en (cit. on pp. 67, 103, 105).

[52] Open-Meteo. *Historical Weather API documentation*. Accessed 2026-02-14. URL: https://open-meteo.com/en/docs/historical-weather-api (cit. on pp. 67, 105).

[53] Open-Meteo. *Solar radiation documentation (timestamping and averaging; GTI tilt/azimuth notes)*. Accessed 2026-02-14. URL: https://open-meteo.com/en/docs (cit. on pp. 67, 103, 105).

[54] Tianqi Chen and Carlos Guestrin. "XGBoost: A Scalable Tree Boosting System". In: *arXiv* (2016). DOI: 10.48550/arXiv.1603.02754. URL: https://arxiv.org/abs/1603.02754 (cit. on p. 67).

[55] Tianqi Chen and Carlos Guestrin. "XGBoost: A Scalable Tree Boosting System". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016, pp. 785–794 (cit. on p. 105).