

Reversing the Flow

Inverse modelling of V2G systems for congestion mitigation

N.J.A. Steinmetz



Reversing the Flow

Inverse modelling of V2G systems for
congestion mitigation

by

N.J.A. Steinmetz

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on 31 August 2026.

Student number:	5186455
Project duration:	February 2026 – August 2026
Thesis committee:	Dr. ir. E. J. L. Chappin, TU Delft, chair
	Dr. ir. I. Nikolic, TU Delft
	L. Schubotz, MSc, TU Delft, PhD supervisor
	J. van Dijk, MSc, EVConsult, company coach

Cover: V2G Hub (Modified)

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.



Preface

Completing this thesis has been a turbulent process, with both rewarding and difficult moments, but above all it has been a period in which I learned a great deal. My work was often driven by boundless curiosity and a desire to take on every possible challenge. I am grateful to everyone who supported me throughout it.

First, I would like to thank my supervisors. Lukas, thank you for your weekly feedback, endless enthusiasm, and for always finding a way forward whenever I arrived with another mountain of problems. Emile, thank you for supervising both my bachelor's and master's theses, for your confidence in me, and for always seeming to know exactly what kind of feedback or direction I needed. Igor, thank you for helping me manage an often fast-moving and wide-ranging thought process, and for your sharp ideas in interpreting the results.

I am also very grateful to EVConsult for giving me the opportunity to conduct my graduation research there and for making the office such an enjoyable place to work. In particular, I would like to thank Jolijn for her positive energy, practical insights, and many helpful ideas throughout the project.

To the friends with whom I shared my studies: thank you for the countless projects, long study evenings in Australia, and all the moments that made getting through this degree both possible and enjoyable.

I am equally grateful to my family, whose unconditional support has carried me through every stage of my studies and continues to inspire me. Above all, I want to express my gratitude to my parents, who have supported me throughout my entire academic journey and have done everything in their power to make every opportunity possible.

Finally, I would like to thank Bart. Through every high and low, you were always there to encourage me, cheer me up, and make everything feel better.

*N.J.A. Steinmetz
Delft, August 2026*

AI Statement

Generative artificial intelligence (AI), primarily ChatGPT using GPT-5.5 and GPT-5.6, was used as a supporting tool throughout the development of this thesis. Its use focused on programming assistance, language editing, document structuring, and as a discussion partner for exploring alternative ways to analyse, interpret, and present the research. AI-generated suggestions were critically reviewed, adapted, and verified before being incorporated. Methodological choices, interpretation of results, and the final content of the thesis remained the responsibility of the author.

For programming and modelling tasks, AI was used to support the development and modification of NetLogo and Python code, to understand existing model logic, and to improve the clarity and efficiency of model and analysis code. For example, AI was asked to help translate an explicit net-aware charging pattern from a source into a model function controlling charging speed. It was also used to develop data-processing and visualisation functions. One example was: “Help me create a function in which I can specify a congestion threshold and penalty exponent and plot median loss against MAD, with optional logarithmic scaling for each axis and support for an exceedance-count loss.” AI also assisted with debugging, adapting scripts for DelftBlue runs, and diagnosing errors encountered during simulation and analysis. Suggested code was subsequently tested and modified where necessary.

AI was also used as a language and document-editing aid. This included shortening and clarifying paragraphs, generating alternative formulations when particular relationships or arguments were not sufficiently clear, translating author-written Dutch text into academic English, converting existing material into LaTeX tables, and helping formulate connecting text based on material already present elsewhere in the thesis. Example requests included: “Can you formulate [2-5 sentences] more clearly and concisely?” and “Elements X and Y do not come through clearly enough in this paragraph; how could this be reformulated?” AI was occasionally used to identify potentially relevant literature or bibliographic information; any sources used in the thesis were independently checked before citation.

Finally, AI was used as a discussion partner for methodological, analytical, and structural choices. This included comparing alternative analysis or presentation approaches against criteria defined by the author, critically reflecting on figures and observed patterns, discussing the logical ordering of sections, and considering alternative interpretations of findings. AI was therefore used to challenge, clarify, and organise the author’s reasoning, rather than to make methodological or interpretative decisions independently. All resulting interpretations and conclusions were evaluated against the underlying model results and, where relevant, the cited literature.

Summary

Vehicle-to-Grid (V2G) can provide flexibility to electricity systems, but its practical contribution to congestion mitigation depends not only on technical battery capacity, but also on user participation, vehicle availability, charging behaviour, and operational control. Existing research often treats participation as fixed or evaluates user acceptance separately from system performance. This study investigates which V2G system configurations are associated with lower local congestion pressure when participation emerges endogenously from heterogeneous contract preferences and behavioural feedback.

A full-year agent-based model of a stylised Dutch urban area was used to simulate the mobility, charging, contract acceptance, and V2G behaviour of 500 electric-vehicle drivers at five-minute resolution. Contract characteristics, charging policies, charging infrastructure, and renewable-energy conditions were varied across a large discrete design space. Congestion pressure was represented by a squared loss function penalising positive net-load exceedances above 600 kW. A learning-guided inverse-modelling procedure, based on a Bayesian Ridge surrogate and two independently seeded learning streams, directed simulations towards potentially favourable configurations. Each proposed configuration was evaluated through repeated stochastic simulations, and Patient Rule Induction Method was used to identify interpretable low-loss design regions.

The results show that participation is an enabling, but insufficient, condition for congestion relief. Higher participation was generally associated with lower loss, with the largest improvements occurring between low and medium-high participation levels and diminishing returns emerging at high participation. Guaranteed range was the strongest predictor of participation, followed by minimum plug-in time, while financial compensation and battery-degradation assumptions had smaller effects. Low congestion loss was most consistently associated with smart charging at residential locations, high renewable-energy coverage, and a medium-to-high guaranteed range. Smart charging and V2G performed complementary functions: controlled charging prevented avoidable EV demand from reinforcing existing peaks, while V2G reduced the remaining positive net load. Nevertheless, congestion remained concentrated in winter evenings, and a limited number of severe peaks persisted.

The inverse search identified favourable configurations relatively early and subsequently improved mainly by avoiding poor regions rather than converging towards a unique optimum. This study contributes by linking heterogeneous contract participation to realised, temporally available grid flexibility and by demonstrating how inverse modelling can identify multiple plausible socio-technical design regions. The findings indicate that effective V2G systems require the joint design of user contracts, charging control, renewable-energy conditions, and grid-responsive operation, rather than the maximisation of participation alone.

Contents

Preface	i
AI Statement	ii
Summary	iii
Nomenclature	viii
1 Introduction	1
1.1 Research problem	1
1.2 Literature gap	2
1.3 Research objectives	2
1.4 Research questions	3
1.5 Research approach	3
1.6 Thesis outline	4
2 Theoretical and empirical background	5
2.1 Congestion	5
2.2 V2G within the EV charging system	7
2.3 V2G research	8
2.3.1 Overview of V2G research	8
2.3.2 Behavioural foundations	8
2.3.3 User preferences in stated-choice studies	9
2.3.4 Related research	11
2.4 Inverse modelling	12
2.4.1 Inverse modelling in agent-based modelling	12
2.4.2 Positioning relative to related modelling methods	13
2.4.3 Surrogate-guided search of parameter spaces	14
2.4.4 Limitations	15
2.4.5 Positioning of the present study	15
2.5 Conceptual synthesis	15
3 Methodology	17
3.1 Modelling approach	17
3.1.1 Agent-based modelling approach	17
3.1.2 Inverse problem formulation	17
3.1.3 Loss function	18
3.1.4 Bayesian Ridge-guided sampling	19
3.2 Experimental setup	20
3.2.1 Experimental structure and terminology	20
3.2.2 Exploratory and final experiment phases	21
3.2.3 Exploratory and final parameter space	21
3.2.4 Repetitions and random-seed design	22
3.2.5 Learning rounds and stopping criterion	23
3.3 Computational implementation	23
3.3.1 NetLogo implementation	23
3.3.2 Python controller and sampler implementation	24
3.3.3 Parallel execution and batch processing	24
3.3.4 Data storage and computational environment	25
3.4 Model verification and validation	26
3.4.1 Verification approach	26

3.4.2	Validation approach	27
3.5	Analysis approach	27
3.5.1	Data preparation and aggregation	27
3.5.2	Exploratory assessment	27
3.5.3	Assessment of sampler learning	28
3.5.4	Parameter and mechanism analysis	28
4	Model description	30
4.1	Model purpose and scope	30
4.2	Model foundations and overview	30
4.3	Model adaptations	31
4.4	Spatial environment and model entities	32
4.4.1	Time and environment	32
4.4.2	Agents	33
4.5	Charging behaviour	33
4.5.1	Non-V2G charging	35
4.5.2	V2G charging and discharging	35
4.6	Behavioural feedback and participation	35
4.7	Energy balance and grid indicators	36
4.8	Data sources and inputs	37
4.9	Model assumptions and simplifications	37
5	Results	39
5.1	Model verification and validation results	39
5.1.1	Verification results	39
5.1.2	Validation results	39
5.2	Exploratory results	39
5.2.1	Run Example	40
5.2.2	Loss function	40
5.2.3	Learning behaviour of the sampler	42
5.2.4	Repetitions, stochasticity and stability	43
5.2.5	Outcome selection and exploratory output relationships	44
5.2.6	Contract and control settings selected for further analysis	46
5.3	Final results	47
5.3.1	Learning behaviour of the sampler	47
5.3.2	Temporal composition of congestion loss	49
5.3.3	Parameters associated with congestion loss	50
5.3.4	Regions associated with low congestion loss	52
5.3.5	Participation and charging-related outcomes	54
5.3.6	Alternative analyses and method selection	56
6	Discussion	58
6.1	Interpretation of results	58
6.1.1	Drivers of V2G participation	58
6.1.2	From participation to usable flexibility	60
6.1.3	System conditions for congestion relief	62
6.1.4	Temporal patterns and residual congestion	64
6.2	Model scope, uncertainty and generalisability	65
6.3	Interpretation, value and limitations of the approach	66
6.3.1	Search performance and surrogate learning	66
6.3.2	Consequences for inference and robustness	66
6.3.3	Positioning and applicability of the approach	68
6.4	Recommendations for future research	69
7	Conclusion	71
7.1	Main research question	71
7.2	Contractual and psychological determinants of participation	72
7.3	Participation and congestion relief	72

7.4	Smart charging, grid-aware charging, and V2G	73
7.5	System-design conditions and the origin of congestion	73
7.6	Methodological contribution of inverse modelling	74
7.7	Implications for system design and policy	74
7.7.1	Aggregators and contract design	74
7.7.2	Grid operation and charging-system	75
7.7.3	Public policy and system planning	75
7.8	Recommendations for future research	75
7.9	Final conclusion	76
References		77
A Evidence on V2G participation preferences and contract design		83
A.1	Behavioural foundations of participation	83
A.2	Contract attributes and experimental levels	84
A.3	Study-specific findings and preference heterogeneity	86
B Experimental design and computational settings		89
B.1	Exploratory and final parameter spaces	89
B.2	Fixed parameter settings	91
B.3	Computational environment and final learning settings	92
C Model		93
C.1	Inherited model foundations	93
C.1.1	Original model by Van der Kam et al.	93
C.1.2	V2G extension by Bakhuis et al.	93
C.2	Agent variables	94
C.3	Detailed charging behaviour	97
C.3.1	Charging access and execution order	97
C.3.2	Non-V2G charging	97
C.3.3	V2G control	98
C.4	Detailed behavioural feedback and participation	100
C.4.1	Latent preference classes and contract evaluation	102
C.4.2	Update timing and inertia	102
C.4.3	Range anxiety	102
C.4.4	Plug-in time concern	102
C.4.5	Environmental self-identity	103
C.4.6	Peer influence	103
C.4.7	Preference updating and contract re-evaluation	103
D Model verification and validation		104
D.1	Verification diagnostics	104
D.2	Supporting validation comparisons	106
D.2.1	Location-specific charging profiles	106
D.2.2	Mobility outcomes	106
D.2.3	Smart-charging and V2G profiles	106
D.2.4	Expert review	107
E Exploratory results		108
E.1	Runs	108
E.2	Loss function	109
E.2.1	Overview tested losses	109
E.2.2	Rank-correlation loss	109
E.2.3	Stability loss	109
E.3	Exploratory learning	111
E.3.1	Stochastic variability of the exceedance-count loss	111
F Final results		114
F.1	Sampler learning	115
F.1.1	Between-round learning	115

F.1.2	Within-round learning	115
F.2	Feature scores rounds	115
F.3	Parameter repetition	118
F.4	Seasonal performance comparison	118
F.5	PRIM	118

Nomenclature

Abbreviations

Abbreviation	Definition
ABM	Agent-Based Model
BD	Battery Degradation
ESI	Environmental Self Identity
EV	Electric Vehicle
FC	Financial Compensation
GR	Guaranteed Range
MP	Minimum Plug-In Time
MAD	Median absolute deviation
PIC	Plug-In Concern
RA	Range Anxiety
SOC	State-of-charge
V2G	Vehicle-To-Grid

Introduction

1.1. Research problem

Climate change requires a rapid transition towards the electrification of energy consumption across sectors such as transport, heating, and industry (International Energy Agency, 2023a). As part of this transition, the rapid growth of electric vehicles (EVs) and the increasing share of intermittent renewable energy sources are placing significant pressure on electricity grids (International Energy Agency, 2023b). The electrification of mobility increases electricity demand, while renewable sources such as wind and solar introduce greater variability and more local energy supply. Together, these developments are challenging the operational limits of existing electricity infrastructure.

In many regions, including the Netherlands, local distribution networks are approaching their capacity limits, leading to increasing instances of grid congestion (Netbeheer Nederland, 2023). Although congestion occurs in both withdrawal and feed-in directions, current Dutch capacity maps and waiting-list data indicate particularly widespread constraints on the demand side (Netbeheer Nederland, 2026; Partners in Energie, 2026). This congestion is becoming a critical barrier for further electrification, as limited grid capacity can delay the integration of new renewable generation and electrified technologies, and in some cases force the curtailment of renewable energy when the network cannot accommodate additional electricity flows (International Energy Agency, 2023b).

Yet, electric vehicles represent a largely untapped source of flexibility within the energy system. Since EVs are parked for the majority of the day, their batteries could potentially be used to support the electricity grid through smart charging strategies and Vehicle-to-Grid (V2G) technologies (Kempton & Tomić, 2005). These technologies enable EVs not only to consume electricity but also to temporarily store energy and feed it back into the grid when needed, helping to balance supply and demand and reduce pressure on the grid. In addition to V2G, various charging strategies such as grid-aware charging and smart charging could help shift or limit charging demand to moments when grid capacity is available (Refa et al., 2023). If coordinated effectively, this flexibility could support more efficient use of existing grid capacity, create additional room for electrification and renewable energy integration, and potentially postpone part of the need for grid reinforcement (International Energy Agency, 2023b).

While the technological development of V2G is rapidly progressing and pilot projects are emerging, successful large-scale implementation depends not only on technical feasibility but also on user participation (Burger, 2025; V2G Hub, 2026). EV owners must be willing to participate in such systems, which introduces behavioural and psychological considerations (Gschwendtner et al., 2021). Participation conditions are commonly formalised through V2G contracts, which can determine, for example, the compensation offered to users or the battery range guaranteed for mobility (Vey et al., 2025). At the same time, high participation alone does not guarantee effective congestion mitigation; the system must be designed in a way that ensures participation translates into grid effectiveness (Yun et al., 2025; Zhang et al., 2024).

Existing research has extensively analysed the technical potential of V2G and the effects of different

charging strategies (Eltohamy et al., 2025). However, these studies typically evaluate predefined scenarios or assume fixed participation levels, providing limited insight into the range of possible system configurations and how they lead to specific outcomes (Sagaria et al., 2024a). As a result, it remains unclear which combinations of design choices are most effective in achieving desired system-level outcomes under realistic behavioural conditions. V2G implementation is therefore not only a technical challenge, but also a socio-technical design challenge in which user behaviour and system operation are closely connected. The interaction between contractual, infrastructural, behavioural, and operational choices creates a large configuration space, making it difficult to determine through predefined scenarios which combinations can produce a desired system outcome.

To address this limitation, this research adopts an inverse-modelling approach. Rather than starting from a predefined system configuration and asking what outcome it produces, inverse modelling starts from a desired outcome and searches the configuration space for combinations capable of producing it. In this research, the target is future-oriented: reducing grid congestion pressure. Instead of focusing on a single best-performing configuration, the approach aims to identify multiple configurations associated with favourable performance. Focusing on Dutch urban areas, the study investigates which V2G configurations can mitigate local grid congestion when participation emerges endogenously from agent behaviour, thereby providing policymakers and system designers with insight into how system design and participation interact to shape grid effectiveness.

1.2. Literature gap

Existing V2G research predominantly focuses on technical potential, often treating behavioural participation as exogenous, fixed, or overly simplified (Gschwendtner et al., 2021; Liu et al., 2025). Studies that do examine user acceptance frequently rely on stated-choice data, stopping at the willingness to participate rather than analysing how this participation dynamically affects actual system performance over time (Geske & Schumann, 2018).

Recent advancements have introduced endogenous participation and heterogeneous preferences, but critical limitations remain. Existing models often rely on short simulation horizons, focus on theoretical rather than actual availability, or omit key behavioural mechanisms such as range anxiety, perceived loss of freedom, and peer influence (Park Lee, 2019; Yun et al., 2025; Zhang et al., 2024). Furthermore, V2G is rarely analysed alongside complementary strategies like smart and grid-aware charging, despite their shared impact on local load patterns.

Consequently, there is a distinct gap for modelling approaches that synthesize realistic mobility, dynamic behavioural responses, and long-term system outcomes. Because the combination of these factors creates a vast and interacting design space, a structured search approach is necessary. Inverse modelling may provide such an approach, but it has predominantly been applied to the reproduction of historical observations (Jensen & Chappin, 2017; Ommering, 2024; Schubotz et al., 2025). Its potential as a forward-looking approach for searching large parameter spaces for configurations associated with future system objectives remains underexplored. This study examines this potential using Bayesian-guided sampling as a structured search method within a stochastic agent-based model.

1.3. Research objectives

The objective of this research is to investigate which Vehicle-to-Grid (V2G) system configurations can contribute to mitigating local grid congestion under endogenous user participation. The study examines how contractual, infrastructural, and operational design choices influence participation and congestion outcomes, and how V2G interacts with and compares to complementary strategies such as smart charging and grid-aware charging.

A secondary methodological objective is to examine how Bayesian-guided inverse modelling supports the search for promising configurations within the large parameter space of a stochastic agent-based model. Such models make exhaustive exploration computationally infeasible, creating a need for search methods that use previous simulations to guide subsequent evaluations towards potentially promising regions. This includes analysing how the Bayesian sampler directs the search and reflecting on the strengths and limitations encountered when applying the approach.

1.4. Research questions

Which Vehicle-to-Grid (V2G) system configurations mitigate local grid congestion in Dutch urban areas when user participation is modelled endogenously?

Sub-questions:

- How do psychological and contractual factors influence EV owners' willingness to participate in V2G systems?
- How does user participation affect the effectiveness of V2G systems in mitigating local grid congestion?
- How do smart charging and grid-aware charging contribute to V2G in mitigating grid congestion under different system design configurations?
- Which combinations of contractual, infrastructural, and operational parameters define favourable congestion outcomes?
- How does Bayesian-guided inverse modelling support the search for promising V2G configurations within the large parameter space of a stochastic agent-based model?

1.5. Research approach

The research approach consists of three linked methodological elements, as illustrated in Figure 1.1. First, the agent-based model simulates mobility, charging behaviour, participation decisions, and resulting grid-load patterns for a candidate system configuration. Second, the inverse-modelling objective evaluates the simulated grid-load pattern by translating the frequency and severity of positive exceedances above a specified congestion threshold into a loss value. Third, a Bayesian-guided sampler learns from previously evaluated configurations and proposes new candidates in more promising parts of the parameter space. Repeating these steps directs the search towards combinations of system parameters associated with this favourable congestion pattern.




Figure 1.1: Conceptual overview of the inverse-modelling framework linking candidate system configurations, agent-based model (ABM), congestion evaluation, and Bayesian-guided sampling.

The first sub-question, concerning the influence of psychological and contractual factors on EV owners' willingness to participate, is addressed through a review of existing literature. This provides the behavioural foundation for modelling participation within the agent-based framework.

The second to fourth sub-questions are addressed using the simulation model in combination with an inverse-modelling loop. By applying Bayesian-guided inverse modelling to search the parameter space defined by key design parameters, including contract characteristics and charging strategies, the study identifies configurations associated with effective congestion mitigation. In particular, the analysis examines how participation affects system effectiveness, how V2G performs in combination

with smart and grid-aware charging, and which system configurations can achieve favourable congestion outcomes.

The fifth sub-question is addressed by analysing the learning behaviour of the sampler, the development of the search across rounds, and the strengths and limitations encountered in applying the approach.

Together, this research design provides a systematic framework to analyse how system design choices translate into system-level performance. As shown in Figure 1.2, the study moves from problem definition and behavioural foundation to agent-based model development, inverse modelling, analysis, and conclusions.




Figure 1.2: Conceptual overview of the research design.

1.6. Thesis outline

The remainder of this thesis is structured as follows. Chapter 2 presents the theoretical and empirical background on grid congestion, V2G, user participation, behavioural theory, and inverse modelling. Chapter 3 describes the methodology, including the modelling strategy, inverse-modelling procedure, loss function, and validation strategy. Chapter 4 details the model structure, agents, behavioural mechanisms, charging strategies, data sources, and assumptions. Chapter 5 presents the verification, validation, exploratory analysis, and final results. Chapter 6 discusses the findings and methodological limitations. Finally, Chapter 7 answers the research questions, outlines implications for system design and policy, and concludes the thesis.

Theoretical and empirical background

This chapter presents the theoretical and empirical foundations of the study. It first discusses local grid congestion and positions V2G within the broader EV charging system, including smart and grid-aware charging. It then reviews existing V2G research, with particular attention to user preferences, behavioural mechanisms, stated-choice evidence, and related participation studies. Finally, it introduces inverse modelling and synthesises the literature into the conceptual basis for the agent-based model and research approach used in this thesis.

2.1. Congestion

Electricity grid congestion occurs when more electricity is requested from the network than it can safely transport at a specific place and time. Although often described as a “full grid”, congestion can occur at different levels of the electricity system. At the transmission level, congestion concerns high-voltage electricity transport between regions. At the distribution level, it concerns the medium- and low-voltage networks that connect households, businesses, local renewable generation, and charging infrastructure.

This thesis focuses on congestion from a local net-load perspective. The model does not represent detailed physical power flows in a specific distribution grid, but simulates how household and service-sector demand, renewable generation, EV charging, and V2G discharging shape aggregate local peak load. This perspective is relevant because EV charging occurs close to homes, workplaces, and public charging locations and can therefore add to local demand peaks. Conversely, smart charging and V2G can reduce or shift EV-related load and may help relieve peak pressure.

Congestion has become increasingly important because the energy transition changes both the volume and timing of electricity flows. Electrification of transport, heating, and industry increases electricity demand, while solar PV and wind generation introduce more variable and decentralised production. These developments make electricity flows less predictable and more locally concentrated. The International Energy Agency argues that electricity grids risk becoming a bottleneck for clean-energy transitions if grid development does not keep pace with electrification and renewable-energy deployment (International Energy Agency, 2023b). In the Netherlands, this challenge is visible in growing waiting lists for transport capacity and connection requests, increasing pressure on grid operators and delaying the electrification or expansion plans of businesses and public organisations. Netbeheer Nederland’s long-term energy-system scenarios show that the transition towards a climate-neutral energy system requires substantial expansion of electricity infrastructure and flexibility resources (Netbeheer Nederland, 2023, 2025).

The electrification of road transport is expected to increase electricity demand further. In March 2026, the Netherlands had approximately 701,000 fully electric passenger cars, representing 7.5% of the fleet, alongside around 58,500 electric commercial vehicles (Rijksdienst voor Ondernemend Nederland, 2026b). These numbers are expected to increase further, as Dutch policy aims for all new passenger cars entering the market in 2030 to be zero-emission, while vans and trucks are also increasingly

electrified (Rijksdienst voor Ondernemend Nederland, 2025; Rijksoverheid, n.d.).

For congestion, charging power is more relevant than annual energy use. Home chargers commonly provide 3.7 or 11 kW, while average Dutch household electricity use corresponds to only about 0.28 kW over a full year (Milieu Centraal, n.d.-a, n.d.-b). An 11 kW charger therefore draws roughly forty times the average household power while charging at peak power, illustrating how simultaneous EV charging can create substantial local demand peaks.

Several complementary solutions are being pursued. The most structural is grid expansion through additional cables, transformers, substations, and high-voltage connections. Such projects have long lead times and can be delayed by spatial constraints, permitting procedures, shortages of staff and materials, and nitrogen-related restrictions. TenneT reports that around two-thirds of the projects intended to address grid congestion are delayed because of the nitrogen crisis (TenneT, n.d.). Grid expansion therefore remains essential but cannot resolve congestion quickly enough on its own.

In addition to grid expansion, congestion can be addressed by making better use of the capacity that is already available. This can be done on both the supply and demand side. Generation may be reduced temporarily when the network cannot accommodate all available output, potentially requiring renewable-energy producers to curtail generation or delay projects (International Energy Agency, 2023b). Electricity demand can instead be lowered, shifted in time, or spread more evenly across the day. Other measures include congestion management, flexible or non-firm connection agreements, time-dependent network tariffs, storage, and conversion of electricity into other energy carriers (Autoriteit Consument & Markt, 2025; CE Delft & Merosch, 2024). These measures do not replace grid reinforcement, but can reduce pressure on the network while new infrastructure is being developed.

Flexibility plays a central role within this portfolio because many forms of electricity use do not need to occur at one exact moment. Industrial processes or EV charging can sometimes be delayed, or operated at lower power without removing the underlying service. This can reduce simultaneous peak demand and better align consumption with variable solar and wind generation. Flexibility can be activated through direct control, agreements with grid operators or aggregators, or financial incentives. Congestion-management contracts compensate users for adjusting consumption or generation when local network capacity is scarce, while flexible transport agreements provide access to the grid under the condition that usage may be restricted during congested periods (Autoriteit Consument & Markt, 2025; CE Delft & Merosch, 2024).

Price signals are another way to influence when electricity is used. Dynamic electricity prices encourage consumption during low-price periods, while time-dependent network tariffs can discourage consumption when network capacity is scarce. However, wholesale prices reflect system-wide supply and demand, whereas congestion is often local, meaning that different price signals do not necessarily produce grid-supportive behaviour. A low market price may therefore coincide with a local network constraint. Price-based flexibility only contributes reliably to congestion reduction when the signal reflects the timing and location of the bottleneck (CE Delft, 2025; Netbeheer Nederland, 2023).

Storage can strengthen these forms of flexibility by moving electricity between periods of surplus and shortage. Batteries can charge when renewable electricity or network capacity is available and discharge during high-demand periods, thereby reducing demand peaks and limiting the curtailment of renewable generation. Its effect nevertheless depends on where the storage is located and how it is controlled. Batteries operated only for energy trading may charge or discharge at unsuitable moments and can even worsen local congestion (CE Delft, 2023; Rijksdienst voor Ondernemend Nederland, 2026a). Effective congestion mitigation therefore requires coordination between grid expansion, flexible demand and generation, storage, pricing, and contractual control.

EV charging is directly connected to this broader solution portfolio. Uncontrolled charging can add to local peak demand, particularly when many vehicles charge after returning home in the evening. EV users and charge-point operators may also face limits on charging power or delays in expanding charging infrastructure where grid capacity is scarce. At the same time, EVs can provide flexibility because charging is often not needed immediately and vehicles are parked for substantial parts of the day. This creates opportunities for charging strategies that better align EV demand with grid conditions. The next section therefore positions Vehicle-to-Grid (V2G) within the broader EV charging system, alongside smart charging and grid-aware charging.

2.2. V2G within the EV charging system

Electric vehicles affect the electricity system through the timing, location, and power level of charging. Under uncontrolled charging, vehicles generally start charging when plugged in and continue until the battery is full or the session ends. Although convenient for users, this can increase local peaks when many vehicles charge during already high-demand periods. Controlled charging instead treats EV demand as flexible, allowing it to be shifted, reduced, or, under bidirectional charging, temporarily reversed.

A first form of flexibility is smart charging. Smart charging refers to strategies in which the timing or power level of EV charging is adjusted using signals and information such as electricity prices, renewable-energy availability, grid loading, and expected departure time (Huber et al., 2019). Rather than reducing total electricity consumption, smart charging distributes demand more favourably over time, for example by delaying charging beyond an evening peak or shifting it towards periods of high renewable generation. It can therefore reduce peak demand and improve renewable-energy use without discharging electricity back to the grid. Smart charging is increasingly used in the Netherlands, particularly at private charging points. In 2025, 41% of surveyed home-charging respondents used dynamic electricity tariffs, while another 32% expressed interest in doing so (ElaadNL & Vereniging Elektrische Rijders, 2025). However, this form of charging is mainly price-driven rather than directly coordinated with local grid conditions.

Grid-aware charging takes a more direct and locally targeted approach. Instead of mainly optimising charging for electricity prices, renewable availability, or user preferences, charging power is adjusted to the available capacity of the local network. This may involve slowing or delaying charging during congested periods through grid-operator signals, contractual agreements, or automatic control. In the Dutch context, this is often referred to as *netbewust laden*. By the winter of 2025/2026, grid-aware charging had been introduced at approximately 5,200 of the roughly 60,000 Dutch public charging points, primarily in areas with high peak loading (Nationale Agenda Laadinfrastructuur, 2025). Grid-aware charging is therefore directly connected to congestion mitigation: it does not return electricity to the grid, but can prevent EV charging from worsening peak-load periods.

Vehicle-to-Grid extends controlled charging by enabling bidirectional electricity flows between EVs and the grid. In a V2G system, a connected EV can charge from the grid when electricity is abundant or demand is low, and discharge electricity back to the grid when demand is high or flexibility is valuable. Kempton and Tomić (2005) define V2G as the use of electric-drive vehicles to provide power to the grid while the vehicle is parked and connected. Because vehicles remain parked for substantial periods, connected and coordinated EV batteries can together function as a distributed flexibility resource.

Although V2G deployment remains at an early stage, the technology is moving beyond demonstration projects towards initial commercial deployment. Numerous pilots have been conducted across Europe and internationally, involving a growing number of bidirectional chargers and participating vehicles (V2G Hub, 2026). More recently, the first commercial V2G propositions have begun to emerge, including offers that compensate users for vehicle availability or discharged electricity and integrated charging or energy contracts (Burger, 2025). These developments indicate a transition from primarily experimental applications towards early market implementation, although large-scale deployment remains limited.

V2G forms part of the broader concept of Vehicle-to-Everything (V2X). Vehicle-to-Load (V2L) supplies individual appliances, while Vehicle-to-Home (V2H) and Vehicle-to-Building (V2B) supply electricity to a home or building, for example as backup power or to increase solar self-consumption. V2G differs by exporting electricity to the wider grid and therefore requires coordination with grid actors and compliance with grid requirements (International Energy Agency, 2026). These applications share the principle of bidirectional charging but differ in the destination of the discharged electricity and the service provided.

V2G can support the electricity system through peak shaving, short-term balancing, and renewable-energy integration. Discharging during high net-load periods can reduce peaks that would otherwise approach or exceed local network capacity is sufficient, while charging during renewable surpluses and discharging during later shortages can improve the use of variable generation. Studies in microgrid and renewable-energy contexts show that controlled charging can increase local renewable-energy use and reduce peak demand, while bidirectional charging provides additional flexibility by allowing EVs to act as temporary storage (van der Kam & van Sark, 2015).

However, the flexibility potential of EVs is not automatically available. Whether an EV can support the grid depends on whether it is available and connected to a suitable charger, and whether sufficient battery energy remains above the driver's mobility needs. V2G also requires bidirectional charging technology, communication between the vehicle, charger, aggregator, and grid operator, and control strategies that determine when charging or discharging should occur. In practice, an aggregator may coordinate many individual EVs and translate their combined flexibility into a grid service, while respecting user constraints such as minimum state of charge, expected departure time, and contract conditions.

V2G is therefore a socio-technical flexibility mechanism whose effectiveness depends on the interaction between infrastructure, mobility behaviour, user participation, contract design, and system operation. Whereas smart and grid-aware charging shift or limit EV demand, V2G can additionally return electricity to the grid. The central question is therefore under which system configurations the resulting participation and additional flexibility contribute to effective congestion mitigation.

2.3. V2G research

The V2G literature can be grouped into three broad streams. The first stream studies the technical, economic, and grid-service potential of V2G, often assuming fixed participation or predefined availability. The second stream studies user acceptance and contract preferences, usually through surveys, interviews, stated-preference experiments, or stated-choice models. The third, smaller stream attempts to link participation to system-level outcomes through simulation or scenario analysis. This section reviews these streams in order to identify what is known about V2G effectiveness, what is known about user participation, and where the connection between both remains underdeveloped.

2.3.1. Overview of V2G research

V2G research has increased substantially in recent years. A large share of this literature examines technical performance, economic and market design, regulatory frameworks, and policy implications (Eltohamy et al., 2025; Gschwendtner et al., 2021). These studies often assume full participation or impose a fixed participation rate (Andersen & Powell, 2025; Huda et al., 2020; Mozafar et al., 2017). In other cases, participation is varied only as a sensitivity parameter to assess system robustness (Bibak & Tekiner-Mogulkoc, 2021; Y. Li et al., 2024; Sagaria et al., 2024a, 2024b). While these approaches provide valuable insights into technical feasibility and economic potential, participation is typically treated as exogenous rather than as a behavioural outcome.

Empirical evidence indicates that participation is an important determinant of V2G system performance and economic viability. Sagaria et al. (2024a) find that V2G can continue to provide valuable grid support when acceptance and battery availability are reduced by up to 30%, indicating a degree of resilience under lower participation. However, the study also shows that higher acceptance improves system reliability and local self-sufficiency. Similarly, Y. Li et al. (2024) reports that increasing participation raises annual revenues in electricity-trading systems. Broneske and Wozabal (2017) emphasises that insufficient user engagement can undermine the financial viability of V2G business models and that contracts must attract users capable of providing valuable flexibility. Taken together, these findings indicate that participation is not merely a background assumption but an important determinant of available flexibility, system reliability, and economic performance.

In parallel with this system-oriented research, another stream of studies focuses on users' behaviour, preferences, attitudes, heterogeneity, and perceived risks regarding V2G participation. Although V2G fundamentally depends on voluntary participation, this behavioural dimension remains comparatively underexplored (Liu et al., 2025; Zhang et al., 2026). This body of work typically relies on interviews, surveys, or stated-preference experiments among current or prospective EV users. Findings suggest that willingness to participate is influenced by multiple factors beyond economic incentives (Gschwendtner et al., 2021), including range anxiety, concerns about battery degradation, financial compensation, and trust in aggregators or system operators (Zhang et al., 2026).

2.3.2. Behavioural foundations

A systematic review of 87 studies by Liu et al. (2025) shows that V2G adoption is shaped by three broad categories of determinants: product characteristics, such as battery degradation and charging flexibil-

ity; individual perceptions, including range anxiety and perceived risk; and economic incentives. The review also identifies a broad range of theoretical approaches, with demand-response models being the most common. These models are useful for representing price responsiveness and demand-side flexibility, but generally provide less detail on how user preferences, perceptions, and social influences shape participation. Despite growing recognition of these behavioural factors, the literature therefore remains strongly focused on technical system design and optimisation, with comparatively limited attention to how users evaluate V2G participation and how their preferences and experiences influence participation decisions.

Behavioural frameworks approach this decision from different perspectives. Technology-acceptance and social-psychological models emphasise perceived usefulness, risk, trust, user control, attitudes, and social influence, while economic and stated-preference approaches focus more directly on the trade-offs users make between participation conditions. A more detailed review of these behavioural frameworks and their application to V2G is provided in Appendix A.1. Of particular relevance to this thesis are random-utility approaches, which represent participation as a choice between alternatives based on the utility associated with their attributes (Huang et al., 2021; Zonneveld, 2019). The model used in this thesis builds on this utility-based decision structure but extends it with dynamic psychological feedback. This allows heterogeneous contract preferences to be represented explicitly while allowing participation to evolve through behavioural factors such as range anxiety, plug-in concern, and environmental self-identity, providing a direct connection to the stated-choice evidence reviewed in the next section.

2.3.3. User preferences in stated-choice studies

User preferences for V2G participation are commonly examined through stated-choice experiments, in which respondents choose between hypothetical contracts with different conditions. Because large-scale V2G participation is still limited and operational data from emerging commercial services are not always publicly available, opportunities to observe actual user choices across different contract designs remain scarce. By varying attributes such as compensation, mobility guarantees, plug-in requirements, battery degradation, and user control, these studies reveal the trade-offs users make when considering participation. Appendix Table A.1 summarises the selected studies, their national contexts, and the attribute levels examined, while Appendix Table A.2 presents their study-specific findings, reported preference heterogeneity, and relevant qualifications. The following subsections synthesise the main cross-study findings for each contract attribute.

Minimum range and SOC

Minimum State of Charge (SOC) and driving-range guarantees consistently emerge as important determinants of V2G participation, reflecting users' need to preserve sufficient energy for mobility. Their importance is closely related to range anxiety and the perceived risk that the vehicle may not be available when needed. Across several studies, mobility guarantees rank among the most influential contract attributes, particularly under normal charging conditions and among user groups primarily concerned with vehicle availability (Bakhuis et al., 2025; Geske & Schumann, 2018; Huang et al., 2021; Meijssen, 2019; Parsons et al., 2014).

The evidence also suggests that responses are not necessarily linear. Moderate reductions in guaranteed battery levels may have little effect, whereas acceptance can decline sharply once the guarantee falls below a level considered sufficient for mobility (Kubli et al., 2018; Vey et al., 2025). The importance of these guarantees may decrease when fast charging is available, as users can restore mobility more quickly (Huang et al., 2021; Meijssen, 2019). Overall, mobility security appears to be a central condition for participation rather than merely one contract benefit among many.

Compensation

The influence of financial compensation is less consistent across studies. In some experiments, remuneration is a primary driver of contract choice, whereas in others it is secondary to mobility guarantees or flexibility requirements (Bakhuis et al., 2025; Geske & Schumann, 2018; Kubli et al., 2018; Meijssen, 2019; Parsons et al., 2014; Zonneveld, 2019). This variation indicates that compensation does not have the same importance for all users or under all contract conditions.

Several studies also find limited sensitivity to moderate differences in variable remuneration, particularly when additional per-kWh payments are offered on top of fixed compensation (Huang et al., 2021; Meijssen, 2019; Vey et al., 2025). However, very low payments can substantially reduce contract utility, suggesting that remuneration may partly operate as a minimum acceptable condition rather than producing proportional increases in participation at every higher level. Compensation therefore matters, but may be insufficient to offset perceived mobility risks, restrictive plug-in requirements, or loss of control.

Flexibility

Flexibility-related attributes determine how strongly participation constrains the everyday use of an EV. Requirements concerning plug-in duration, connection timing, the number of required plug-in days, and the possibility to opt out can therefore have a substantial influence on contract preferences. Several studies identify groups for whom these conditions are among the most important attributes, demonstrating considerable heterogeneity in the value users place on maintaining control over when their vehicle is connected (Bakhuis et al., 2025; Geske & Schumann, 2018; Meijssen, 2019; Vey et al., 2025; Yun et al., 2025).

The importance of plug-in requirements appears to depend on the charging context and the levels included in the experiment. Plug-in time was more influential under fast charging than under normal charging, whereas a study using relatively short connection durations found no significant effect (Meijssen, 2019; Zonneveld, 2019). Highly flexible arrangements, including frequent opt-out opportunities, are generally valued. These findings suggest that the availability of vehicles for V2G cannot be treated only as a technical condition, because stricter availability requirements may themselves reduce willingness to participate.

Battery degradation

Battery degradation is generally perceived as a cost or risk of V2G participation, and higher numbers of discharge cycles tend to reduce contract acceptance. Nevertheless, degradation rarely emerges as the dominant contract attribute and is often less influential than mobility guarantees or plug-in flexibility (Bakhuis et al., 2025; Huang et al., 2021; Meijssen, 2019).

The evidence should be interpreted cautiously because respondents may find degradation expressed in numbers of cycles difficult to understand. This may explain the counterintuitive, non-monotonic relationship between battery degradation and utility reported by Zonneveld (2019). Qualitative evidence nevertheless confirms that expected battery lifetime remains an important concern, particularly among private EV owners (Hu et al., 2025). Battery degradation is therefore relevant, but its measured importance depends partly on how clearly the consequences are communicated.

Contract duration

Evidence on contract duration is mixed. Some studies find no statistically significant effect, while others report preferences for either shorter or longer contracts (Huang et al., 2021; Kubli et al., 2018; Meijssen, 2019; Zonneveld, 2019). These differences suggest that contract length is evaluated in relation to the wider contract design rather than in isolation. Shorter contracts may reduce perceived commitment, whereas longer contracts may be valued when they offer stability and avoid repeated renegotiation.

The value attached to daily opt-out possibilities further indicates that users may be more concerned with retaining flexibility and control than with contract duration itself (Vey et al., 2025). Contract length is therefore unlikely to have one uniform effect; its influence depends on whether users perceive the arrangement as stable, restrictive, or easily reversible.

Overall, the findings show that V2G contract preferences are shaped by trade-offs between compensation, mobility security, flexibility, battery use, and contractual commitment. Mobility guarantees and control over plug-in behaviour are valued relatively consistently, whereas the influence of compensation and contract duration varies more strongly between studies and user groups. Financial compensation therefore appears to have a threshold effect. Very low payments can reduce willingness to participate, but the literature does not consistently indicate that high compensation is necessary when mobility guarantees, flexibility, and user control are adequately addressed.

Heterogeneity

Preferences for V2G participation are highly heterogeneous. Some users primarily value financial compensation, whereas others place greater importance on mobility security, plug-in flexibility, or control over vehicle use (Bakhuis et al., 2025). This means that participation cannot be represented adequately by a single average preference profile and that identical contract conditions may produce different responses across user groups.

Socio-demographic characteristics such as income, age, and gender explain part of this variation, but their effects are inconsistent across studies. Mobility-related concerns, particularly range anxiety and the desire to retain control over vehicle availability, appear to distinguish user groups more consistently than demographics or environmental attitudes (Bakhuis et al., 2025; Geske & Schumann, 2018; Parsons et al., 2014; Vey et al., 2025). Overall, the evidence suggests that V2G contracts should accommodate distinct preference profiles rather than assume uniform participation across the EV population. Detailed study-specific findings and qualifications are provided in Appendix Table A.2.

2.3.4. Related research

Despite these advances, relatively limited attention is paid to how heterogeneous acceptance outcomes translate into system-level effectiveness over time. Although a small number of studies attempt to model participation endogenously, their scope remains constrained.

(Park Lee, 2019) model different V2G contract types within a microgrid system that includes hydrogen storage and photovoltaic (PV) generation. Participation is treated endogenously in the sense that contract structures influence household involvement. The model simulates 200 houses and 50 drivers over monthly runs across one year. However, the participation assumptions are not empirically grounded in observed behavioural data. Moreover, while drivers are included conceptually, the model primarily focuses on household-level interactions rather than dynamic driving behaviour. Although this study integrates contract variation into system modelling, the absence of empirically grounded behavioural parameters and feedback-factors such as range anxiety or environmental self-identity limits the realism of the participation outcomes.

Zhang et al. (2024) advance the field by modeling acceptance probability as a function of compensation, perception of charging rules, and perceived freedom of travel. The study claims to be the first to endogenously model participation within a large-scale agent-based framework, including 30,000 agents. Heterogeneity is represented through distributions of user characteristics, and the acceptance model includes attributes related to compensation, free-travel rules, and perceived charging-rule constraints. However, the model is constrained to a single-day simulation and is not dynamic over time. There is no explicit modeling of driving behaviour, and no integration of renewable energy generation. The primary output is hourly available capacity, which does not necessarily translate into effective system contribution relative to energy demand or renewable integration. While the scale of the model is impressive, the temporal limitation and lack of explicit mobility behaviour reduce its ability to capture longer-term participation dynamics or cumulative system effects.

Yun et al. (2025) takes a different approach by combining survey data with a latent class model to account for heterogeneity in user preferences. The study varies contract attributes such as monetary compensation, required connection days, distance to charger, weekday connection time, and guaranteed SOC. System-level impacts are assessed using hourly generation and demand predictions, and contracts are modeled as predefined scenarios. This top-down calculation framework allows for estimation of peak shaving and net load reduction under different participation assumptions. However, it is not an agent-based model and participation is evaluated within static scenario structures rather than as a dynamic process. Similar to Zhang et al. (2024), the analysis is limited to short-term horizons, which restricts insight into adaptation, learning, or evolving acceptance over time.

Across these studies, participation is treated endogenously, representing a clear advancement compared to models assuming fixed participation rates. Nevertheless, several limitations persist. First, behavioural effects are often simplified or restricted to contract preferences rather than integrated with explicit mobility behaviour. Second, most models operate on short time horizons, frequently limited to a single day, which prevents the observation of dynamic effects, learning processes, or long-term reliability implications. Third, system effectiveness is commonly measured in terms of hourly capacity or

peak reduction, without fully assessing energy balance, renewable integration, or cumulative economic performance over time.

As a result, while these studies begin to address the relationship between participation and system effectiveness, they do not yet provide a fully integrated dynamic framework in which heterogeneous, empirically grounded acceptance decisions interact with mobility behaviour and long-term system outcomes. This gap highlights the need for modeling approaches that simultaneously capture endogenous participation, behavioural heterogeneity, and dynamic system effectiveness over extended time horizons.

Because V2G does not yet have a dominant institutional or contractual design, the relevant research problem is not only to evaluate one predefined V2G scenario, but to explore which combinations of contract conditions, charging strategies, and behavioural responses can produce acceptable congestion outcomes. This creates a link to inverse modelling, which starts from a desired system outcome and searches for the conditions that can generate it.

2.4. Inverse modelling

Inverse modelling is a modelling approach that starts from an outcome of interest and works backwards to identify the conditions, assumptions, parameter configurations, or model structures that could generate that outcome (Epstein, 2023). This contrasts with conventional forward modelling, where a model starts from predefined assumptions and simulates the outcomes that follow from them. In simple terms, forward modelling asks what happens if certain conditions are imposed, while inverse modelling asks which conditions could lead to a desired or observed result. This section first defines inverse modelling, then discusses related methods, and its main limitations.

2.4.1. Inverse modelling in agent-based modelling

In agent-based modelling, inverse modelling changes the organising direction of conventional forward simulation: rather than specifying behavioural rules, parameter values, and system assumptions and examining the outcomes that follow, it starts from a target outcome and searches for inputs or model formulations that can reproduce or achieve it.

Inverse modelling is not one uniformly defined computational method. Inverse problems have a long methodological tradition in applied mathematics and fields including engineering, geophysics, and medical imaging, where they commonly involve estimating unknown parameters or inputs from observed outputs (Yaman et al., 2013). Their application to agent-based modelling and generative social science is more recent (Epstein, 2023). Within this literature, inverse approaches differ in the target used, the model elements being inferred, and the method used to conduct the search.

The target of an inverse process may be synthetic or stylised, such as flocking or segregation (Epstein, 2023; Gunaratne et al., 2023), or empirical, such as an observed time series (Vu et al., 2023). To guide the search, it must be translated into a quantitative criterion that measures how closely each simulated outcome corresponds to the target and indicates which candidate configurations are most promising.

The elements being inferred may range from numerical parameter values (Ommering, 2024) to behavioural rules and model procedures (Epstein, 2023; Schubotz et al., 2025), and alternative model structures (Jensen & Chappin, 2017; Schubotz et al., 2025). Inverse applications may therefore serve different purposes, including parameter calibration, identifying plausible behavioural mechanisms, comparing alternative model structures, and developing or refining behavioural theory. The corresponding search can be conducted using different computational methods, including random and Bayesian search (Ommering, 2024), genetic algorithms (Schubotz et al., 2025), and genetic programming or related evolutionary approaches (Epstein, 2023).

One influential development in agent-based modelling is Inverse Generative Social Science (iGSS). Whereas conventional generative social science examines whether specified micro-level rules produce a macro-level pattern, iGSS starts from such a pattern and searches for rules that can generate it (Epstein, 2023). These rules may be constructed from behavioural components and logical or mathematical operators rather than limited to parameter adjustments. The aim is to identify interpretable and plausible generative explanations, not only formulations with a close numerical fit. Since different rules

can produce similar outcomes, candidate explanations must also be assessed for theoretical plausibility and their ability to reproduce patterns not used during the search.

Related inverse approaches restrict the structural search space more strongly. Jensen and Chappin (2017) present an automated modelling procedure for innovation diffusion in which predefined decision-model variants, structural alternatives, and parameter values are evaluated against empirical diffusion patterns. A matching function combines these patterns and measures how well each candidate model reproduces them. The approach thereby combines automated model selection and parameter optimisation to identify model variants that can serve as plausible explanations of historical diffusion. Unlike iGSS, it searches within a bounded, modeller-defined library rather than constructing new agent rules from elementary building blocks.

Schubotz et al. (2025) develop a related structural approach referred to as structural sculpting. Their procedure changes the presence and ordering of existing model procedures and evaluates the resulting structures against historical solar PV adoption data. Unlike more open forms of iGSS, the search does not construct completely new behavioural functions from elementary operators. Instead, it recombines components that already exist within the model. This makes it possible to investigate whether different structural formulations can reproduce the same empirical behaviour and whether particular procedures appear necessary, substitutable, or redundant.

A more parameter-oriented form of inverse modelling is used by Ommering (2024). In this application, the structure of an existing solar PV adoption model remains largely fixed, while behavioural utility weights are varied using Random Search and Bayesian Search to reproduce municipal adoption curves. This approach lies close to conventional parameter calibration because it searches for parameter values that improve correspondence with historical observations. It can nevertheless be understood as parameter-based inverse modelling because it works backwards from observed outputs to behavioural parameter configurations that are consistent with the observed development. The resulting interpretations remain conditional on the assumed model structure.

Although these approaches differ in implementation, they share a common analytical purpose: reducing the set of model configurations that could plausibly generate a specified outcome. Depending on the application, this may involve identifying behavioural parameter values, selecting between alternative model structures, or constructing new agent rules. The resulting configurations do not necessarily reveal one true explanation, but define behavioural or structural hypotheses that remain consistent with the target.

The applications discussed above use inverse modelling primarily for explanatory purposes. Empirical applications seek configurations that reproduce observed historical behaviour, while synthetic and stylised applications test which micro-level rules can generate a predefined system-level phenomenon. Although not all targets are therefore historical, they generally represent behaviour that is to be reproduced and explained rather than a future performance objective that is to be achieved. The inverse logic may also be used for the latter purpose by starting from a desired system outcome and searching for the conditions capable of producing it. This use can be understood as a design-oriented or future-oriented form of inverse modelling, in which the target is not an observed pattern to be explained but a prospective system state to be realised. To position this broader use of inverse logic, it is useful to compare it with related modelling approaches that may employ similar computational procedures.

2.4.2. Positioning relative to related modelling methods

Inverse modelling overlaps with several established modelling approaches. The boundaries are not strict, since similar computational procedures may be described differently depending on the target, the elements being varied, and the interpretation of the results. The distinctions made here therefore concern the organising research question rather than the exclusive use of particular algorithms.

Calibration and parameter estimation are closely related to parameter-based inverse modelling. Calibration searches for parameter values that improve the correspondence between simulated outcomes and empirical observations (Fagiolo et al., 2007). When the model structure is assumed to be fixed, the same search may be understood either as calibration or as a parameter-based inverse problem. The distinction depends mainly on whether the resulting values are treated as fitted model parameters or are used to reason about plausible underlying conditions or mechanisms.

Calibration may also extend beyond numerical parameters. Structural calibration compares alternative behavioural rules, procedures, or model formulations based on their ability to reproduce empirical evidence. It therefore overlaps strongly with structural forms of inverse modelling, although the emphasis remains on improving the empirical adequacy of the model. The present research differs from both conventional and structural calibration because it does not seek to estimate the parameters or structure of an existing historical V2G system.

Pattern-oriented modelling is closely related to explanatory inverse modelling. It deliberately uses multiple observed patterns, often at different organisational, spatial, or temporal scales, to constrain plausible parameter values, behavioural rules, or model structures (Grimm et al., 2005). A formulation gains credibility when it can reproduce several independent patterns simultaneously. The present research is not pattern-oriented in this strict sense. Compared with automated structural inverse approaches such as iGSS, however, standard pattern-oriented modelling typically evaluates a relatively small set of candidate model structures specified by the modeller, rather than systematically constructing or evolving structures through an automated search process (Vu et al., 2023). Its search is guided by one central performance measure, congestion-related positive net-load exceedance, while other outcomes are used as plausibility and side-effect checks rather than as simultaneous empirical constraints.

Inverse modelling may be implemented using optimisation procedures, but in explanatory inverse applications optimisation functions primarily as a search mechanism. Candidate parameters, rules, or structures are not evaluated only to identify the best-performing model, but are interpreted as hypotheses about the mechanisms capable of generating the target behaviour. Optimisation, by contrast, is primarily concerned with identifying one set of configurations that perform favourably according to a specified objective.

Exploratory modelling generally begins with a broad range of uncertain inputs and decision alternatives and examines the outcomes produced across an ensemble of forward simulations (Bankes, 1993; Kwakkel et al., 2016; Lempert et al., 2003). Its purpose is often to map uncertainty, compare policies, and assess robustness across multiple plausible futures. However, exploratory-modelling frameworks may also use directed search to identify input combinations associated with outcomes of interest. When the model structure is largely predefined, such procedures can be computationally very similar to a performance-guided inverse search.

2.4.3. Surrogate-guided search of parameter spaces

Operationalising inverse modelling requires systematic exploration of the relevant parameter or design space. Parameter combinations may be selected randomly, exhaustively evaluated, or sampled using a predefined non-adaptive experimental design. Such approaches do not learn from previous simulation outcomes and can therefore become inefficient when the parameter space is large and individual model runs are computationally expensive (K. Li et al., 2026).

Alternatively, the search can be guided by information from previously evaluated configurations. A learning algorithm estimates the relationship between model inputs and the objective value and uses this estimate to determine which configurations should be evaluated next. Such methods can be used both to reproduce observed outputs and to search for configurations that achieve a desired system outcome.

Random search can perform relatively well in high-dimensional spaces despite its simplicity (Bergstra & Bengio, 2012). It does not, however, use the results of previous evaluations to guide later samples. Bayesian optimisation addresses this limitation by constructing a probabilistic surrogate model of the objective function and using its predictions and uncertainty to select promising new evaluations (Snoek et al., 2012).

Within such learning-guided approaches, candidate selection involves a trade-off between exploitation and exploration. Exploitation prioritises configurations predicted to perform well, whereas exploration prioritises uncertain parts of the parameter space. Intermediate strategies combine both objectives. The effectiveness of such a search depends not only on the surrogate model, but also on how candidate configurations are generated and ranked.

Bayesian Ridge regression offers a relatively lightweight probabilistic surrogate. It treats regression

coefficients as random variables and estimates regularisation within a Bayesian framework (MacKay, 1992; scikit-learn API, 2026). Although its linear structure limits its ability to represent complex non-linear response surfaces, it can provide both predictions and uncertainty estimates at relatively low computational cost.

2.4.4. Limitations

Several limitations follow from the inverse approach. In empirical inverse modelling, for example when plausible behavioural mechanisms are inferred by searching for model formulations that reproduce an observed historical pattern, a first limitation is the availability and suitability of data that can serve as the target. When candidate models are evaluated against only one dataset or a limited set of patterns, there is also a risk of overfitting to the selected target: a close fit may reflect dataset-specific noise or compensating model assumptions rather than the underlying system dynamics (Epstein, 2023; Renzini et al., 2023). Independent observations or additional patterns are therefore important for assessing whether an inferred mechanism generalises beyond the data used to guide the search.

In the present performance-oriented application, the analogous risk is not overfitting to empirical observations, but over-specialising the search to the modelled objective. A parameter combination may achieve a favourable numerical score without producing desirable behaviour on system dimensions that are omitted from the objective.

The outcomes remain conditional on the model structure, parameter ranges, input data, representation of future system conditions, behavioural assumptions and their formalisation, as well as on the objective used to evaluate system performance. Together, these choices define which configurations can be explored, which model outcomes can occur, and which parts of the parameter space the search treats as promising. In particular, the loss function does not merely evaluate the resulting simulations but actively directs subsequent sampling towards configurations that perform well according to the selected definition of congestion. The resulting promising regions are therefore conditional on both the model formulation and the chosen objective, and the analysis cannot rule out relevant mechanisms, designs, or alternative favourable regions that would emerge under different assumptions or performance criteria.

Finally, inverse modelling inherits the stochastic variability and complex, potentially nonlinear dynamics of the underlying model, which can make relationships between parameters and system-level outcomes difficult to identify consistently (Platt, 2020).

2.4.5. Positioning of the present study

The present research combines elements of parameter-based inverse modelling and directed exploratory modelling. It is inverse in the sense that a specified system-level performance guides the search towards input combinations capable of producing it. At the same time, it is exploratory because the search covers uncertain future conditions and a broad range of system configurations rather than estimating one true parameter set or reconstructing one historical mechanism.

The search space includes numerical parameters, uncertain energy-system and participation conditions, categorical design choices, and predefined behavioural alternatives. Charging-policy choices activate different implemented behavioural rules, so the analysis includes limited structural variation in the active behavioural logic of the model. It does not, however, generate new agent rules or restructure the underlying model architecture. The approach is therefore best described as a performance-guided inverse exploration of a predefined but behaviourally explicit ABM, aimed at identifying promising regions of the V2G system space.

2.5. Conceptual synthesis

This chapter has positioned V2G as a socio-technical flexibility mechanism within the broader EV charging system. Grid congestion creates a need for flexibility, but the extent to which EVs can contribute depends on more than technical battery capacity. Smart charging and grid-aware charging can shift or limit EV demand, while V2G can additionally provide electricity back to the grid during peak-pressure periods. However, this potential only becomes usable when vehicles are available, connected, sufficiently charged, and when users are willing to participate under acceptable contract conditions.

The literature shows that V2G participation is shaped by a combination of contractual, behavioural, and operational factors. Contract attributes such as financial compensation, guaranteed range, minimum plug-in time, and battery degradation influence willingness to participate, while psychological factors such as range anxiety, environmental self-identity, plug-in concern, and peer influence may affect how users evaluate participation over time. At the same time, system-level effectiveness depends on how this participation translates into actual charging and discharging behaviour under mobility and infrastructure constraints.

This creates the central modelling challenge for this thesis: V2G effectiveness cannot be understood by analysing technical capacity or user acceptance in isolation. Instead, participation, charging behaviour, and local net-load outcomes must be analysed together. The resulting design space is large and multidimensional, as contractual, behavioural, infrastructural, and operational conditions interact and may produce similar system-level outcomes.

Therefore, this research combines an agent-based model representing heterogeneous EV users with a performance-guided inverse exploration of a predefined system space. Rather than reproducing an observed historical pattern or inferring one true parameter set, the search starts from a desired congestion-related outcome and identifies multiple configurations capable of producing favourable performance. A Bayesian Ridge-guided sampling procedure directs computational effort towards promising parts of the large configuration space, after which the resulting data is analysed to identify recurring conditions and favourable regions. The approach is therefore future- and design-oriented, while its conclusions remain conditional on the model structure, explored parameter space, and selected performance objective.

3

Methodology

The overall research design and its relation to the research questions were introduced in Chapter 1. This chapter explains how that design was operationalised. It first describes the agent-based and inverse-modelling approaches, followed by the experimental setup and computational implementation. The chapter then presents the procedures used for model verification, validation, and analysis.

3.1. Modelling approach

3.1.1. Agent-based modelling approach

This study uses an agent-based modelling (ABM) approach to analyse how participation in V2G systems influences system-level outcomes. The model builds on an existing EV-charging and V2G participation model, which is extended and applied here to analyse congestion-oriented charging and discharging behaviour.

The approach is appropriate for this study for three reasons. First, as highlighted by Zhang et al. (2024), agent-based modelling can capture differences in user preferences, decision processes, and behavioural responses. This is essential for modelling participation as an endogenous outcome rather than a fixed assumption. Second, ABM captures how system-level outcomes, such as grid congestion, emerge from decentralised decisions by heterogeneous actors. Third, its time-based structure enables mobility patterns, charging behaviour, and electricity-system dynamics to be represented simultaneously, including their interactions and stochastic variation over time.

3.1.2. Inverse problem formulation

The inverse-modelling procedure is applied as a targeted future-oriented search for V2G configurations associated with a pattern of lower local peak-load exceedance. Whereas the forward model maps a parameter configuration to a simulated net-load trajectory, the inverse procedure starts from the desired system outcome and searches for configurations that produce it. The procedure starts from a congestion-related performance objective of low congestion pressure. This objective is operationalised through a loss function that penalises the frequency and severity of positive exceedances of a predefined net-load threshold. Each parameter configuration is evaluated by running the agent-based model and calculating the corresponding loss value from the resulting net-load trajectory. Configurations with lower loss values are interpreted as more effective at reducing local congestion pressure under the assumptions of the model. The conceptual basis and positioning of inverse modelling are discussed in Section 2.4.

The aim is not to derive a single optimal or definitive solution, but to examine whether the procedure can guide the search towards favourable regions of the parameter space and to analyse the conditions under which these configurations perform well. An inverse formulation is suitable for this purpose because the study starts from desired future system behaviour and works backwards to identify multiple combinations of behavioural, contractual, and technical conditions capable of producing it.

Because the parameter space is large and each stochastic model configuration must be evaluated repeatedly, exhaustive enumeration is computationally infeasible. The search is therefore conducted iteratively using a surrogate model that learns from previously evaluated configurations and guides subsequent evaluations towards potentially favourable parts of the parameter space. The loss function and sampling strategy are described below, while the actual parameter space and experiment settings are specified in Section 3.2.

3.1.3. Loss function

The inverse-modelling procedure evaluates each candidate configuration using a congestion-related loss function. The loss serves both as a measure of simulated congestion performance and as the performance signal that guides the sampler. The choice of this function is central to the analysis because it defines what constitutes good performance and determines how the sampler prioritises different regions of the parameter space. Each inverse-modelling run is guided by one predefined loss specification.

For a simulation run, the loss $\mathcal{L}_{\theta,p}$, for threshold θ and penalty exponent p , is defined as the cumulative normalised exceedance above θ , penalised according to exponent p :

$$\mathcal{L}_{\theta,p} = \sum_{t=1}^T \left(\frac{\max(N_t - \theta, 0)}{\theta} \right)^p, \quad (3.1)$$

where N_t is the simulated local net load at time step t , θ is the congestion threshold, p is the exceedance penalty exponent, and T is the number of simulated time steps. Time steps below the threshold receive no penalty. Above the threshold, the penalty increases with the magnitude of the exceedance, with higher values of p assigning relatively more weight to severe peaks.

Dividing the exceedance by θ normalises it relative to the threshold against which it is evaluated. This makes the penalty dimensionless and allows losses based on different thresholds to be compared on a more consistent relative scale.

The penalties are then summed over all time steps, so that the loss captures the cumulative frequency and severity of threshold exceedance over the simulation period. Within each experiment series, all compared runs use the same simulation horizon and time-step duration. Summing rather than averaging therefore does not affect their relative ranking; it only changes the numerical scale of the loss. The summed formulation was retained to express the total exceedance burden over the simulated period and to avoid unnecessarily small loss values.

Only upper net-load exceedance is included in the loss. Negative net load is interpreted as a renewable-energy surplus that may require curtailment, but curtailment and lower network constraints fall outside the congestion objective of this study. Participation, charging and discharging volumes, battery state of charge, and low-battery outcomes are also excluded. These outcomes are analysed separately to distinguish congestion performance from the behavioural and operational mechanisms through which it is achieved. The loss should therefore be interpreted specifically as an upper-peak congestion objective and not as a composite measure of overall system performance.

The threshold and penalty exponent determine which part of the net-load profile is emphasised and may therefore affect the configurations preferred by the sampler. Lower thresholds penalise a broader portion of the net-load profile and increasingly shift the objective from exceptional peak reduction towards more general load reduction, whereas higher thresholds focus more narrowly on extreme peaks but may produce many zero or near-zero losses. Similarly, low penalty exponents retain sensitivity to a relatively broad set of exceedance events, whereas high exponents concentrate the loss increasingly on a small number of severe peaks. The exploratory comparison therefore sought a specification that remained informative across configurations while still assigning greater importance to operationally relevant high-load periods.

Before selecting the final loss specification, alternative congestion thresholds and penalty forms were compared. Thresholds between 400 and 800 kW were examined, with particular attention to the 500–600 kW range and additional higher thresholds to test formulations focused more narrowly on extreme

peaks. The tested penalty forms included an exceedance count, a linear penalty, and polynomial penalties with $p = 2, 3, 10$, and 50 . These alternatives were compared to assess whether they preserved sufficient variation between configurations, produced similar configuration rankings, and yielded comparable relationships with the sampled model parameters.

The complete net-load time series was retained for every simulation, allowing the alternative loss formulations to be calculated retrospectively from the same simulated trajectories. These post-hoc calculations support sensitivity analysis and comparison between loss definitions, but do not alter the original search trajectory: the sampler in each inverse-modelling run learns only from the loss specification used during that run.

The exploratory comparison and the empirical basis for this selection are presented in Section 5.2.2. Based on this assessment, the final learning phase used a congestion threshold of $\theta = 600$ kW and a squared exceedance penalty with $p = 2$. Once selected, this specification was held constant throughout the final learning experiments.

Figure 3.1 illustrates how threshold exceedances are translated into the congestion-loss penalty.




Figure 3.1: Illustration of the congestion-loss calculation for $\theta = 600$ kW and $p = 2$. Panel (a) shows exceedance of the net-load threshold, while Panel (b) shows the corresponding penalty at each time step.

3.1.4. Bayesian Ridge-guided sampling

The principles of surrogate-guided search and Bayesian Ridge regression are discussed in Section 2.4.3. In this study, the learning-guided sampling procedure uses Bayesian Ridge regression as a lightweight surrogate model rather than implementing a complete conventional Bayesian-optimisation framework. Bayesian Ridge is a probabilistic linear regression method in which the regression coefficients are treated as random variables and regularisation is estimated within a Bayesian framework (MacKay, 1992; scikit-learn API, 2026). The procedure is therefore best interpreted as a Bayesian Ridge-guided inverse search. Figure 3.2 presents its iterative workflow.

When no preloaded training data were available, the procedure started with a set of randomly selected configurations. These initial observations formed the training history shown on the left of Figure 3.2. The surrogate was then trained on the encoded parameter settings and their corresponding congestion losses and used to predict the loss and predictive uncertainty of new candidate configurations.

An *experiment series* refers to one complete execution of the Bayesian sampling procedure. Within an experiment series, learning proceeded through successive *Bayesian batches*, with each Bayesian batch representing one sampler-update cycle. At the start of a batch, a new pool of unevaluated can-




Figure 3.2: Workflow of the Bayesian Ridge-guided sampling procedure. Candidate generation is shown in pink, surrogate training and candidate ranking in blue, and agent-based simulation and loss evaluation in orange.

candidate configurations was generated from the predefined parameter space, represented in pink in Figure 3.2. The trained surrogate predicted the loss and uncertainty of these candidates, after which they were ranked according to their acquisition scores. Depending on the sampler setting, selection prioritised favourable predicted loss, low predictive uncertainty, or a weighted combination of both (scikit-learn API, 2026). The configurations with the most favourable scores were selected as the next Bayesian batch.

Each selected configuration was subsequently evaluated through repeated agent-based simulations, shown in orange in Figure 3.2. Once the simulations had been completed, their congestion losses were added to the cumulative training history and the surrogate was refitted before a new candidate pool was generated. The feedback arrow in the figure therefore shows how newly evaluated configurations informed subsequent sampling iterations.

Later experiment series could be initialised using successfully completed observations from preceding experiments. These preloaded observations formed the initial training history, after which each experiment series learned cumulatively from its own newly evaluated configurations.

The surrogate was fitted to the losses of individual stochastic repetitions. Repeated evaluations of the same parameter configuration therefore entered the training data as separate observations rather than being aggregated into a configuration-level mean or median. This retained repetition-level variation in the training data, although the Bayesian Ridge model did not explicitly model the full distribution of stochastic outcomes associated with each configuration. Aggregation across repetitions was applied only during the subsequent assessment of configuration performance and sampler learning.

Because Bayesian Ridge imposes a regularised linear approximation, it was not expected to reproduce the complete nonlinear and stochastic response surface of the agent-based model. Its purpose was instead to provide a computationally inexpensive approximation that was sufficiently informative to rank candidates and guide simulations towards potentially favourable regions. Because the surrogate model, candidate-generation procedure, and acquisition rule simplify the full objective landscape, the procedure was used to improve the direction and efficiency of the search rather than to guarantee identification of a unique global optimum. Its adequacy was assessed empirically through changes in proposal quality, the development of the search across learning rounds, and consistency between parallel experiment series.

3.2. Experimental setup

3.2.1. Experimental structure and terminology

The simulation campaign consisted of several nested experimental levels, illustrated in Figure 3.3. The figure serves as a reference for the terminology used throughout this chapter. Terms corresponding to the levels shown in the figure are italicised when they are first introduced. An individual *run* refers to one complete execution of the NetLogo model. During such a run, the model is initialised, simulated for




Figure 3.3: Overview of the parallel Bayesian sampling structure. The two independent learning chains (V2G 1 and V2G 2) run in parallel and follow the same internal structure. Each chain consists of successive learning rounds, each learning round contains two parallel experiment series (A and B), and each experiment series consists of multiple Bayesian batches. The enlarged panel illustrates one Bayesian batch, in which four parameter combinations are selected and each is evaluated in ten repetitions (runs) using the same parameter settings but different random seeds.

the specified time horizon, and the selected output variables are recorded. A *parameter combination* refers to one unique set of values for the variable model inputs. Because the model is stochastic, the same parameter combination can produce different outcomes when evaluated repeatedly. A *repetition* therefore refers to one run of the same parameter combination using a different NetLogo random seed.

3.2.2. Exploratory and final experiment phases

The simulation campaign consisted of an exploratory phase and a final learning phase. During the exploratory phase, separate experiment series were used to assess the loss formulation, parameter ranges, sampler implementation, simulation horizon, and number of repetitions. Because these experiments addressed different methodological questions and used different settings, they were not treated as one continuous learning sequence.

The final learning phase used the selected congestion-loss formulation, a full-year simulation horizon, the reduced final parameter space, ten repetitions per parameter combination, and the corrected exploitative sampling procedure. All final experiment series used the same model specifications, allowing their results to be compared and combined within the defined learning structure.

Two independently seeded *learning chains*, denoted V2G 1 and V2G 2, were maintained. Each chain consisted of successive *learning rounds*, with two independent experiment series, A and B, executed in parallel within each round. The two series started from the same preloaded training history and sampling settings, but did not exchange newly evaluated observations during execution. Their successfully completed results were combined to construct the preloaded training history for the subsequent learning round. The final experiments used purely exploitative candidate ranking based on predicted loss, predicted uncertainty was not included in the final candidate score.

3.2.3. Exploratory and final parameter space

The exploratory phase examined a broader parameter space to identify settings that produced structurally implausible, highly unstable, or technically uninformative model behaviour. Such settings were excluded from the final search to avoid spending the limited computational budget outside the intended

design domain. During this phase, the congestion-responsive charging thresholds were also varied and subsequently fixed at 450 kW for charge reduction, 700 kW for charging cessation, 550 kW for discharge initiation, and 700 kW for full discharge. Parameters not retained in the final search were held constant at the values reported in Appendix B.2. The complete exploratory parameter space is reported in Appendix B.1.

The final learning phase used the reduced parameter space shown in Table 3.1. It retained variation in contract design, charging infrastructure, renewable-energy conditions, and charging policy, while excluding parameter values and settings that had proved unsuitable or insufficiently informative during the exploratory phase.

The contract attributes and energy-system parameters were based on the V2G model and stated-choice experiment by Bakhuis et al. (2025). The charging-infrastructure sliders were varied by $\pm 20\%$ around their reference values, and all available charging-policy options were evaluated for each charging location.

Table 3.1: Parameter space used during the final learning phase

Category	Parameter	Final values
Energy system	renewable-coverage	40, 50, 60
Energy system	solar-vs-wind	40, 50, 60
Charging infrastructure	private-chargers-slider	0.8, 0.9, 1.0, 1.1, 1.2
Charging infrastructure	public-chargers-slider	0.8, 0.9, 1.0, 1.1, 1.2
Charging infrastructure	office-chargers-slider	0.8, 0.9, 1.0, 1.1, 1.2
Charging infrastructure	commercial-chargers-slider	0.8, 0.9, 1.0, 1.1, 1.2
Charging policy	home-charging-policy	normal, net-aware, smart
Charging policy	work-charging-policy	normal, net-aware, smart
Charging policy	public-charging-policy	normal, net-aware, smart
Contract design	contract-FC	30, 60, 90, 120, 150
Contract design	contract-GR	15, 45, 75, 105, 135
Contract design	contract-MP	1, 3, 5, 7, 9
Contract design	contract-BD	1, 2, 3, 4

The resulting final parameter space contained 75,937,500 possible discrete parameter combinations. Exhaustive enumeration was therefore computationally infeasible, motivating the use of the learning-guided sampling procedure. The complete final parameter space, together with the broader exploratory space, is reported in Appendix B.1.

3.2.4. Repetitions and random-seed design

The exploratory experiments used varying numbers of repetitions. Lower repetition counts in the initial experiments allowed a broader part of the parameter space and a larger number of methodological settings to be examined within the available computational budget. The final learning phase used ten repetitions per parameter configuration, based on the repetition-stability analysis presented in Section 5.2.4, to represent stochastic variation more consistently.

Two separate forms of random seeding were used. NetLogo random seeds controlled stochastic processes within the agent-based model, including variation in mobility and agent behaviour. The Python-wrapper seed controlled the stochastic generation of candidate configurations within the sampling procedure. The wrapper seed was therefore distinct from the NetLogo seeds used to evaluate the selected configurations.

Within an experiment series, the same set of ten NetLogo seeds was reused across parameter configurations. This common-seed design improved comparability by evaluating alternative configurations against the same set of stochastic starting sequences. At the same time, the use of multiple seeds prevented the sampler from learning only from one specific mobility. Seed sets were renewed between learning rounds and differed between the independently seeded learning streams.

Different Python-wrapper seeds were used for the separate learning streams. As a result, the streams

generated and selected different sequences of candidate configurations even when they started from comparable parameter spaces and training histories. Together, the wrapper- and model-level seed design reduced dependence on both one particular search trajectory and one particular set of simulated stochastic conditions. All seeds were recorded with the corresponding experiment settings and outputs to support reproducibility.

3.2.5. Learning rounds and stopping criterion

An operational stopping rule was used to prevent the learning procedure from continuing indefinitely. Repetitions were first aggregated at the level of the parameter combination. For each experiment series, the median loss of each parameter combination was calculated, after which the distribution of these configuration-level median losses was summarised.

Let $m_{r,s}$ denote the median of the configuration-level median losses in round r and learning stream s , and let $u_{r,s}$ denote the corresponding upper quartile. To reduce sensitivity to stochastic variation between individual rounds, all criteria had to be satisfied for two consecutive learning rounds. The learning process was considered to have reached a plateau when:

1. the round-to-round reduction in the average of the two stream-level medians was less than 10%. This criterion indicated that the typical performance of newly evaluated configurations was no longer improving substantially;
2. the reduction in the average upper quartile was less than 10%. This ensured that stabilisation was not limited to the centre of the distribution while the higher-loss configurations were still improving substantially;
3. the difference between the two stream-level medians was less than 25% of their average. This required the parallel streams to show sufficiently similar performance, reducing the likelihood that an apparent plateau was specific to one search trajectory.

This stopping rule should be interpreted as a transparent, study-specific computational criterion rather than as proof that the global optimum of the parameter space had been identified.

3.3. Computational implementation

The computational framework connects the NetLogo agent-based model to an external Python simulation controller, which orchestrates the execution of the model and the inverse-search procedure. NetLogo contains the model agents, behavioural rules, mobility processes, charging and discharging procedures, and electricity-system calculations. The Python controller manages the simulation campaign by loading experiment settings, generating and selecting parameter configurations, assigning random seeds, executing the model, collecting outputs, calculating congestion-related loss, updating the sampling model, and storing completed results.

Figure 3.4 presents the overall computational workflow. Fixed parameters and input data are combined with the variable parameters selected by the sampler and transferred to the NetLogo model. Each selected configuration is then evaluated through repeated simulations using controlled random seeds. Repetition-level losses are returned to the sampler as training observations, whereas outcomes are aggregated across repetitions only for the subsequent analysis of configuration performance and sampler learning.

3.3.1. NetLogo implementation

The agent-based model was implemented in NetLogo (Wilensky & Rand, 2015). As shown in the central part of Figure 3.4, each simulation combined fixed model settings and input data with the variable parameter values selected for the evaluated configuration. The assigned NetLogo random seed was set before model initialisation, after which the model was executed in headless mode for the specified simulation horizon.

Selected NetLogo reporters recorded the time-dependent electricity-system, charging, discharging, participation, mobility, and state-of-charge outcomes required for the subsequent analysis. The internal model procedures and behavioural rules are described in Chapter 4. The present section therefore focuses on the computational execution of the model rather than repeating its conceptual structure.




Figure 3.4: Computational workflow of the surrogate-guided inverse-modelling procedure. Candidate configurations are generated from the defined parameter space and evaluated through repeated NetLogo simulations, and the resulting repetition-level congestion losses are returned to the surrogate for subsequent learning.

3.3.2. Python controller and sampler implementation

Model execution was controlled through a Python wrapper based on PyNetLogo (Jaxa-Rozen & Kwakkel, 2018) and the parallel sampling framework developed by Schubotz (2025). The controller loaded the experiment configuration and parameter space, generated candidate configurations, assigned repetition identifiers and random seeds, and transferred the settings to independent NetLogo instances.

After each simulation, the controller retrieved the requested reporters and calculated the congestion-related loss from the simulated net-load time series. As illustrated in Figure 3.4, the loss was calculated separately for each stochastic repetition and returned to the sampler together with the evaluated parameter settings. The parameter values, repetition identifier, random seed, model outputs, and calculated loss were retained together, allowing every observation to be traced to its experimental settings.

For Bayesian Ridge-guided experiments, the Python controller, also implemented candidate generation, surrogate prediction, candidate selection, preloading of previous training observations, and updating of the surrogate after completed Bayesian batches. The Bayesian sampler was adapted from the approach developed by Schubotz et al. (2026). The candidate-generation and acquisition settings were specified separately for each experiment series, allowing exploratory and final experiments to use different sampler configurations without changing the underlying execution workflow.

The selected configurations followed the feedback loop shown in the upper part of Figure 3.4. After their repetitions had been completed, valid repetition-level losses were added to the training data and used to update the Bayesian sampler before the next candidate pool was ranked.

3.3.3. Parallel execution and batch processing

Parallel execution was used because the inverse-modelling procedure required a large number of independent NetLogo runs. Each worker process operated an independent NetLogo instance in headless mode. Simulation tasks did not exchange information during execution and could therefore be distributed across the available worker processes.

A task consisted of one parameter configuration, one repetition identifier, and one NetLogo random seed. Tasks were assigned to workers using a chunk size of one. This limited the number of tasks reserved by an individual worker and improved load balancing because simulation runtimes could vary

with participation levels, charging-control settings, and other model conditions.

Figure 3.5 illustrates this process for a selected simulation batch. In the final learning setting shown, a candidate pool of 100 configurations was first ranked, after which four configurations were selected. Expanding each selected configuration into ten stochastic repetitions produced 40 independent simulation tasks. These tasks were placed in a shared queue and executed by four independent NetLogo workers as soon as worker capacity became available.




Figure 3.5: Parallel execution of a selected simulation batch. Four configurations are expanded into repetition-level tasks and distributed dynamically through a shared queue across four independent NetLogo workers. Repetition-level outputs and congestion losses are collected until all tasks in the batch have completed.

Within Bayesian experiments, a set of parameter configurations was selected before each surrogate update. All repetitions belonging to the selected configurations were then executed in parallel. The surrogate was updated only after the corresponding Bayesian batch had been completed and its valid loss observations had been returned. This Bayesian batch was distinct from the larger candidate pool evaluated computationally by the surrogate: the candidate pool contained unevaluated alternatives, whereas the Bayesian batch contained the configurations selected for actual NetLogo simulation.

Parallel tasks could finish in a different order from the order in which they were submitted. Each returned result nevertheless remained linked to its parameter values, repetition identifier, and random seed. Parallelisation therefore changed the execution time and completion order but did not alter the conceptual model logic or the relation between model inputs and outputs.

3.3.4. Data storage and computational environment

Results were stored after each completed execution batch rather than only after the complete experiment series had finished. Each stored batch contained the sampled parameter values, repetition identifiers, random seeds, selected model reporters, and calculated congestion-related losses. Intermediate storage reduced the risk of losing completed simulations when an experiment exceeded the available computational wall time or was otherwise interrupted.

An experiment series that reached the HPC wall-time limit was not automatically excluded. A stored batch was retained when the file could be read successfully and contained the expected simulation rows, parameter values, reporters, random seeds, and valid loss values. Incomplete, corrupted, or technically failed outputs were excluded before further learning or analysis. A partially completed experiment series could therefore contribute its successfully completed batches even when the full planned series was not finished.

The main experiments were executed on the DelftBlue high-performance computing environment (DHPC), 2024. NetLogo was operated in headless mode through Java, while Python controlled candidate generation, task distribution, parallel execution, post-processing, learning, and data storage. The software environment, parallelisation, simulation, randomisation, and sampler settings are reported in Appendix B.3.

3.4. Model verification and validation

Verification and validation were conducted to assess whether the model was implemented according to its conceptual design and whether its behaviour was sufficiently credible for the purpose of this study. Verification concerns the correct implementation and integration of model procedures, whereas validation concerns the adequacy of the model as a representation of the relevant real-world system (Dam et al., 2013). Both are required because technically correct code does not necessarily produce credible behaviour, while plausible outputs may conceal implementation errors.

The model was not developed entirely from scratch. It builds on the V2G model developed by Bakhuis (2026), which in turn extends the electric-vehicle charging model of Van Der Kam et al. (2019). The inherited mobility, charging, and energy-system components were therefore not comprehensively revalidated. Verification focused primarily on components introduced or substantially modified in this study and on their interaction with the inherited model structure. These included the V2G controller, contract-based participation, behavioural feedback mechanisms, congestion-related loss calculation, and the Python and NetLogo experimental procedure.

Validation was treated as an assessment of fitness for purpose rather than as proof that the model can predict a future V2G system exactly. Large-scale residential V2G participation and the modelled contract structures cannot yet be compared with a complete empirical reference system. Validation therefore combined temporal and behavioural pattern checks, comparison with theoretical expectations and existing studies, and qualitative expert review. Table 3.2 summarises the verification and validation design.

Table 3.2: Overview of the verification and validation design.

Assessment	Model aspect	Procedure and purpose
Verification	Battery and power constraints	Component and boundary tests of battery state of charge, charging and discharging power, guaranteed-range constraints, simultaneous charging and discharging, and the classification of driven distance as empty-battery, low-battery, or regular electric driving.
Verification	V2G control	Targeted runs with different guaranteed-range values and control thresholds, including comparison of requested and applied charging and discharging setpoints.
Verification	Participation mechanisms	Agent-level tracing of contract utility, latent-class assignment, participation thresholds, and the updating of psychological and social variables.
Verification	Model integration	Full-model tests of net-load accounting, charging-strategy activation, and interactions between vehicles and home, workplace, public, and commercial charging infrastructure.
Verification	Experimental pipeline	Checks of parameter transfer, random-seed assignment, reporter collection, run completeness, loss calculation, and the traceability of stored observations.
Validation	Temporal patterns	Inspection of daily and seasonal mobility, electricity-demand, renewable-generation, vehicle-availability, and net-load patterns.
Validation	Charging behaviour	Comparison of normal and net-aware charging profiles with empirical charging data, and assessment of smart and V2G charging behaviour against theoretical expectations and mobility constraints.
Validation	External credibility	Comparison with findings from relevant V2G and participation studies and qualitative review by an EV-sector expert.

3.4.1. Verification approach

Verification was performed iteratively during model development using component-level tests, targeted boundary and extreme-value runs, agent- and tick-level tracing, and integration tests of the complete model (Dam et al., 2013). Short diagnostic runs were used to inspect battery and power constraints, the sequence of charging and V2G decisions, guaranteed-range protection, contract and behavioural

updates, charging-location allocation, driving-distance classification, and net-load accounting. Longer runs were used to identify inconsistencies that emerged over time.

The Python and NetLogo workflow was verified separately by checking parameter and random-seed transfer, reporter completeness, expected run counts, loss calculation, and the traceability of stored results.

3.4.2. Validation approach

Validation followed a pattern-oriented and fitness-for-purpose approach. The temporal behaviour of the inherited mobility, demand, and renewable-generation components was inspected alongside the charging and discharging behaviour introduced in this study. Normal and net-aware charging profiles were compared qualitatively with empirical location-specific charging data, whereas smart and V2G behaviour were assessed against their intended control logic, mobility constraints, and relevant scenario studies. Qualitative expert review was used to assess the practical plausibility of the charging assumptions and resulting load profiles.

Supporting verification diagnostics and validation comparisons are reported in Appendix D. The main outcomes are presented in Section 5.1.

3.5. Analysis approach

The simulation output was analysed to determine which parameter configurations were associated with lower congestion-related loss, how consistently these configurations performed across stochastic repetitions, and through which behavioural and operational mechanisms these outcomes occurred. The analysis was exploratory: the applied methods were used to identify recurring patterns, influential parameters, and favourable regions of the parameter space, and to characterise the conditions and mechanisms associated with lower loss, rather than to estimate isolated causal effects or identify a uniquely optimal configuration.

3.5.1. Data preparation and aggregation

Completed batch files from the exploratory and final experiments were combined with their associated experiment metadata. Individual simulation runs were retained as separate observations and remained linked to their parameter configuration, repetition identifier, random seed, experiment series, learning stream, and learning round.

Before analysis, the outputs were checked for technical failures, missing reporters, invalid loss values, and incomplete simulation time series. Mobility outcomes, including low-battery and empty-driving events, were used as plausibility indicators rather than as additional components of the congestion-loss function.

For analyses of parameter-configuration performance, repetitions were aggregated at the level of the unique sampled parameter configuration. Because the loss distributions were skewed and could contain extreme outcomes, the median was used as the main measure of central performance. The interquartile range and median absolute deviation (MAD) were used to describe variation across repetitions. Individual repetition-level observations were retained for analyses of stochastic variation and for assessing the training behaviour of the sampler. In addition, a combined measure was used to capture central performance and robustness, calculated as the median loss plus two times the MAD. This combined measure penalises parameter combinations that perform well centrally but are unstable across repetitions.

3.5.2. Exploratory assessment

The exploratory experiments were used to assess the methodological choices that preceded the final learning phase. Alternative congestion-loss formulations were calculated retrospectively from the stored net-load time series and were not fed back to the sampler. They therefore re-evaluated the same simulated configurations without changing the realised search trajectory, although using a different loss during sampling could have resulted in different proposed configurations. The alternative formulations were compared in terms of their numerical distributions, the frequency of zero or near-zero losses, the ranking of parameter configurations, and their relationships with the sampled parameters. Rank cor-

relations were used to assess the extent to which alternative thresholds and penalty forms identified similar configurations as relatively favourable or unfavourable.

Two distinct aspects of stochastic stability were assessed. For the comparison between loss formulations, repetition-to-repetition variability was measured within each parameter configuration using the median absolute deviation (MAD). Because the numerical scale differed strongly between loss definitions, this variability was considered relative both to the typical loss level and to the spread in median performance between parameter configurations. This diagnostic assessed whether a loss formulation amplified stochastic variation relative to the differences it was intended to distinguish. Separately, dedicated experiments with 50 repetitions were used to assess how many repetitions were required to obtain a stable estimate of configuration performance. Estimates based on smaller repetition samples were compared with the corresponding 50-repetition reference for median loss, MAD, and configuration ranking. The exploratory experiments were also used to identify parameter values or ranges that produced technically invalid, structurally implausible, or uninformative model behaviour before defining the final parameter space.

3.5.3. Assessment of sampler learning

Sampler learning was assessed both between successive learning rounds and within individual experiment series. For the between-round analysis, each sampled parameter configuration was first summarised by its median loss across the repetitions. The resulting proposal-level losses from the two parallel experiment series were pooled within each learning chain and round. Their distribution was then summarised using the median and the lower and upper quartiles. Round-to-round changes in these statistics were used to assess whether later rounds increasingly concentrated the sampled configurations in lower-loss regions. Alternative aggregation methods, including the median and trimmed mean across repetitions, were compared as a robustness check.

The predefined stopping criterion was evaluated separately. For each parallel experiment series, the median and upper quartile of proposal-level loss were calculated before the corresponding values of the two series were averaged. Round-to-round changes in these averages were compared with the stopping thresholds. The relative difference between the two series medians was used as an additional measure of consistency between the parallel search trajectories.

Learning within an experiment series was assessed across its successive internal Bayesian batches. The median proposal loss of the first two batches was compared with that of the final two batches. Mean batch loss was inspected as a supporting measure to determine whether apparent improvement was caused mainly by changes in the typical proposal or by the reduction of extreme high-loss proposals. Differences between the two parallel experiment series were also examined to assess the consistency of within-series learning.

3.5.4. Parameter and mechanism analysis

The parameter and mechanism analysis aimed to identify which sampled inputs were most informative for congestion performance, which combinations of parameter values were associated with favourable outcomes, and through which model mechanisms these outcomes emerged. Several complementary post-hoc methods were tested.

The final analytical strategy drew primarily on principles of Exploratory Modelling and Analysis (EMA). EMA examines patterns across ensembles of model experiments rather than relying on a single reference simulation or point prediction (Bankes, 1993). The EMA Workbench developed by Kwakkel (2017) was used to implement the analyses available within this framework. In this study, the ensemble consisted of the configurations evaluated during the exploratory and final experiments. Because the final configurations were selected adaptively by the sampler, the analysis describes patterns within the evaluated configurations rather than their prevalence across the complete parameter space.

Global parameter relevance was assessed primarily using an Extra Trees regression model implemented through the EMA Workbench and fitted as a post-hoc predictive model. Extra Trees combines multiple randomised decision trees and can represent nonlinear relationships and interactions without imposing a predefined functional form. Because the loss distribution was strongly right-skewed, the principal analysis used log-transformed configuration-level median loss. The model's out-of-bag predictive per-

formance was assessed before its feature-importance scores were interpreted, and results obtained using raw and log-transformed loss were compared to evaluate the stability of the identified rankings.

Scenario discovery was subsequently performed using the Patient Rule Induction Method (PRIM) implemented through the EMA Workbench. PRIM searches for multidimensional boxes containing a relatively high concentration of cases that meet a predefined outcome of interest. In this study, binary outcome definitions were used to identify regions associated with outcomes such as low congestion loss and high participation. Candidate boxes were compared using coverage, representing the proportion of all cases of interest captured by a box, and density, representing the proportion of observations within that box that met the outcome definition. The resulting boxes were interpreted as associative descriptions of the evaluated configurations rather than as causal or exhaustive rules for the full parameter space.

Several additional methods were tested as supporting analyses. Permutation importance assessed how predictive performance changed when the information contained in an individual parameter was disrupted, providing a robustness check on the Extra Trees feature ranking. Selected SHAP analyses decomposed tree-model predictions into parameter contributions and were used to examine the direction and variation of effects for individual observations and across the evaluated data. Partial dependence plots summarised the average relationship learned by the predictive model between selected parameters and loss. Descriptive distributions, rank correlations, and direct comparisons of simulation outcomes were used to determine whether the model-based patterns were also visible in the underlying data.

Finally, mechanism-focused analyses related contract attributes to participation and compared charging policies, infrastructure, and V2G operation using charging and discharging volumes, location-specific profiles, net-load trajectories, state of charge, and mobility outcomes. These analyses were used to distinguish whether favourable congestion outcomes resulted from avoided charging demand, shifted charging, V2G discharge, changes in participation, or combinations of these mechanisms.

4

Model description

This chapter describes the agent-based model used to analyse how EV drivers' mobility, charging behaviour, and V2G participation affect local electricity-system outcomes. It explains the model purpose and scope, the main agents and system components, the charging and discharging rules, and the behavioural feedback mechanisms. The chapter ends by outlining the key inputs, outputs, assumptions, and simplifications used in the simulation experiments.

4.1. Model purpose and scope

The purpose of the model is to simulate how electric vehicle drivers interact with charging infrastructure across a stylised urban area, and how their charging, discharging, and participation behaviour affects aggregate electricity demand. The main objective is to evaluate whether Vehicle-to-Grid (V2G) participation and flexible charging strategies can reduce local peak-load exceedance and thereby relieve congestion pressure. In addition, the model is used to analyse related outcomes, including net electricity load, EV charging demand, V2G discharging supply, V2G participation, state-of-charge outcomes, and renewable-energy utilisation.

The model represents a spatially simplified urban system with residential, office, commercial, and public charging locations. A fixed population of EV drivers follows daily travel profiles and moves between home, work, and commercial destinations. Non-V2G charging behaviour can follow normal, net-aware, or smart charging rules, while V2G-enabled drivers may also discharge when connected and when the controller identifies a local system need. Electricity demand from households and the service sector, together with solar and wind production, creates time-varying net-load conditions under which EV charging and discharging take place. Simulations are run for one year, allowing the model to capture repeated daily and seasonal patterns in mobility, demand, and renewable generation.

The scope is limited to the interaction between EV charging behaviour, V2G participation, renewable-energy availability, and aggregate local peak-load outcomes. The model does not aim to reproduce detailed distribution-grid power flows, electricity-market price formation, or physical battery degradation processes. Charging and discharging are represented through simplified power rates per tick, constrained by battery capacity, charging speed, charger availability, mobility needs, and minimum battery requirements. Contract attributes such as financial compensation, guaranteed range, minimum plug-in time, and battery degradation compensation affect whether drivers participate in V2G through the contract evaluation and behavioural feedback mechanisms.

4.2. Model foundations and overview

The model used in this study builds on two successive agent-based models. The EV charging model developed by Van Der Kam et al. (2019) provides the underlying mobility, electricity-system, and charging-infrastructure structure, while the V2G extension by Bakhuis (2026) adds contract evaluation, endogenous participation, behavioural feedback, bidirectional charging, and aggregator control. The latter forms the direct starting point for the present study. A more detailed descriptions of the two underlying

models are provided in the Appendix Section C.1.

Figure 4.1 presents the modular structure of the resulting model and indicates which elements were inherited, modified, or added in the present study.




Figure 4.1: Modular overview of the agent-based V2G model. Line styles and markers indicate whether model elements were inherited, adapted, or added in the present study.

This section describes how the model components interact during a simulation run. System-design parameters, including V2G contract attributes, charging strategies, renewable-energy supply, and infrastructure levels, define the conditions under which each simulation is executed. These parameters affect the mobility, energy-input, and infrastructure modules, which together determine when vehicles are available, how much background demand and renewable production are present, and where charging opportunities exist.

At the centre of the model are EV driver agents. Each driver is characterised by mobility needs, location, charging mode, and V2G participation status. Participation is influenced by the behavioural module, which includes contract evaluation, latent preference classes, range anxiety, plug-in concern, environmental self-identity, and peer influence. These behavioural mechanisms determine whether drivers participate in V2G and how participation preferences evolve over time.

Charging and discharging decisions are handled by the charging module. Non-V2G vehicles can follow normal, net-aware, or smart charging rules, while V2G participants can also be charged or discharged under aggregator control. These charging and discharging actions feed into the energy-balance module, where EV demand and V2G supply are combined with building demand and renewable generation to calculate aggregate net load.

The model outputs include V2G participation, charging and discharging volumes, net-load indicators, state-of-charge and mobility outcomes, and behavioural indicators. Through these modules, the model connects contract design, behavioural participation, mobility, infrastructure, and energy-system conditions to the congestion-related outcomes analysed in this study.

4.3. Model adaptations

The Bakhuis model was adapted to fit a different research objective. Whereas the preceding model focuses on V2G adoption, energy balancing, aggregator strategies, and profitability, the present study focuses on local peak-load reduction under varying system-design conditions. The adaptations therefore mainly concern the experimental representation of contracts, the control of V2G and non-V2G charging, the response of the aggregator to system conditions, and several behavioural and operational refinements.

First, the contract structure was changed from a choice between multiple predefined contracts to a

single offered contract configuration per simulation run. Drivers still evaluate whether to participate based on contract utility, latent preference heterogeneity, and behavioural feedback, but they no longer choose between several competing V2G contracts within the same run. This makes contract attributes experimental design parameters rather than alternatives in an in-model contract menu. The adaptation is necessary for the inverse-modelling setup, because each run evaluates how one specific combination of financial compensation, guaranteed range, minimum plug-in time, and battery degradation compensation affects participation and congestion outcomes.

Second, the charging logic was extended to compare V2G with other congestion-relevant charging strategies. In the Bakhuis model, non-participants charge independently, while V2G participants hand control to the aggregator, the aggregator then coordinates charging and discharging under contractual constraints, especially the guaranteed range. The present model retains this distinction, but introduces location-specific non-V2G charging policies. Normal, net-aware, and smart charging can be assigned separately to residential, work, and commercial charging locations. This allows the model to analyse whether congestion reduction comes from V2G discharging itself, from better timing of non-V2G charging, or from combinations of both.

Third, the aggregator logic was revised from an imbalance-market and profitability-oriented controller to a congestion-oriented controller. The original Bakhuis model includes an aggregator that offers contracts, manages participating vehicles, and earns imbalance-market revenues. In the present study, aggregator profit is removed from the main model logic, because the objective is not to analyse business-case profitability but to evaluate peak-load reduction. The updated aggregator therefore responds to local net-load pressure and available battery flexibility. It prioritises restoring vehicles below their guaranteed range, allows flexible V2G charging when this does not worsen peak pressure, and activates V2G discharging when net load exceeds the congestion-relevant threshold.

Fourth, several behavioural and operational refinements were made to improve consistency with the congestion-focused analysis. Environmental self-identity is updated based on whether realised charging or discharging behaviour is grid-supportive at the moment it occurs. The plug-in concern update was also corrected to ensure that plug-in experiences are evaluated consistently over time. The social charging option from the original model was removed, as it is outside the scope of the present V2G contract-design analysis.

Finally, several implementation corrections and extensions were made. The accounting of empty-driving events was corrected to ensure that mobility outcomes were reported accurately. The charger reservation procedure was revised so that vehicles reserve a charging point only after satisfying the conditions required to charge at that location. Additional reporters were introduced to trace location-specific charging and discharging and to report state-of-charge and mobility outcomes separately for participating and non-participating vehicles.

Together, these adaptations shift the model from a V2G adoption and profitability model towards a congestion-focused simulation environment. The resulting model can be used to analyse how contract design, behavioural participation, charging strategies, infrastructure conditions, and energy-system settings interact to affect V2G availability, charging and discharging behaviour, state-of-charge outcomes, and local peak-load exceedance.

4.4. Spatial environment and model entities

4.4.1. Time and environment

The model represents a stylised urban area divided into residential, office, and commercial zones. The simulation world consists of grid patches representing areas of 2×2 km. This spatial scale was calibrated such that simulated travel distances approximate Dutch travel patterns. EV drivers move between home, workplace, and commercial destinations, where charging opportunities depend on the available charging infrastructure.

Each simulation tick represents five minutes. Each run covers 96,768 ticks, corresponding to twelve four-week months. Detailed grid-cell and environment state variables are reported in Tables C.2 and C.3 in Appendix C.

4.4.2. Agents

The model contains 500 EV-driver agents. Each driver is either a V2G participant or a non-participant and follows day-specific mobility profiles that determine when and where the vehicle travels. Drivers have individual attributes including state of charge, access to a home charger, contract preferences, and psychological characteristics. The main EV-driver state variables are reported in Appendix C.2.

V2G participation is re-evaluated weekly. Each driver belongs to a latent preference class and has individual sensitivities to the contract attributes and psychological factors represented in the model. These characteristics determine the utility of the offered contract and the participation threshold against which it is evaluated. Participation can therefore change over time as drivers accumulate mobility, charging, and V2G experiences.

The model also includes a single aggregator, represented as a central coordinating entity for participating vehicles. It controls their charging and discharging through rule-based responses to the current energy balance and net load, with the objective of limiting local peak-load pressure while respecting vehicle availability, battery constraints, and guaranteed range. The aggregator does not use demand forecasts, formal optimisation, or observed electricity-market prices.

4.5. Charging behaviour

Charging behaviour depends on a vehicle's location, charger availability, state of charge, participation status, and the charging policy assigned to the relevant location. Vehicles can charge or discharge only after reaching their destination and obtaining access to a suitable charging point. Non-V2G and V2G charging are controlled separately but update the same system energy balance. Figure 4.2 summarises the sequence within each simulation tick, from the baseline energy balance and non-V2G charging to emergency charging, flexible V2G charging, discharging, and the final energy balance. The main decision rules are described below, while the complete operational logic and control settings are reported in Appendix C.3.

Sequential charging and energy-balance updates within each simulation tick

The upper row shows the model sequence; the lower row provides a visual explanation of each stage.




Figure 4.2: Sequential charging and energy-balance updates within each simulation tick. Normal, net-aware, and smart charging form alternative non-V2G policies. After their effect has been incorporated into the residual energy balance, V2G emergency charging is applied first, followed by collectively coordinated flexible charging and peak-responsive discharging.

4.5.1. Non-V2G charging

Non-V2G drivers can follow normal, net-aware, or smart charging. The selected policy determines the charging power once a vehicle has reached an available charger and satisfies the conditions for charging at that location. Access to charging away from home also depends on whether the driver has a home charger and whether sufficient battery energy remains for the expected return journey. The three alternative non-V2G policies and their different temporal responses are illustrated in Figure 4.2.

Under normal charging, an eligible vehicle charges at the available charging power until its battery is full or the vehicle leaves the charging location. Net-aware charging uses the same eligibility rules but limits charging power according to a fixed daily profile. Charging power is reduced to a minimum of 4 kW between 17:00 and 23:00 and is gradually restored after the evening peak, thereby shifting part of the load without using real-time system conditions.

Smart charging determines charging power dynamically. The rule aims to reach a target state of charge of 80% before the next expected departure. It first calculates the minimum charging rate required to reach this target and may use part of the remaining charging capacity when the energy balance indicates sufficient grid room. Charging can be delayed during elevated net load when the vehicle can still reach the target after waiting. Smart charging is therefore rule-based and mobility-constrained rather than based on formal optimisation or forecasts. It primarily spreads charging over the available connection period in response to mobility needs and the energy balance, rather than to minimise charging costs.

4.5.2. V2G charging and discharging

Participating vehicles plug in when they have reached their destination and a suitable charging point is available. The aggregator subsequently determines whether each connected V2G vehicle charges, remains idle, or discharges. Its control is based on the current energy balance and net load and is applied sequentially in the three V2G stages shown in Figure 4.2.

First, vehicles below the battery level associated with their guaranteed range receive emergency charging. This takes priority over the system condition and restores the contractual minimum battery level. Second, the remaining connected vehicles may charge flexibly. An internal energy-balance signal permits charging during renewable surplus and limited maintenance charging close to balance, while a soft peak guard progressively restricts charging as net load increases. The available collective charging budget is distributed with greater priority given to vehicles further below the V2G target state of charge.

Third, eligible vehicles may discharge when current net load indicates peak pressure. Only energy above the guaranteed minimum battery level is available, and vehicles that have already charged in the same tick cannot discharge. The collective discharge budget increases with the amount by which net load exceeds the discharge threshold and is distributed according to the usable battery energy of the connected vehicles.

The four V2G peak-control thresholds were varied during the exploratory phase and subsequently fixed for final learning. Other operational settings, including the internal energy-balance signal, response factors, the smart-charging target and peak-avoidance rule, specific charging speeds, and imbalance-threshold were not part of this exploratory search but remained important design settings. Their functions and values are documented in Appendix C.3, while the exploratory ranges and final values of the four varied thresholds are reported in Table B.1.

4.6. Behavioural feedback and participation

The model represents V2G participation as a recurring contract decision that can change through experience. Figure 4.3 summarises this process. Contract attributes and heterogeneous driver preferences determine contract utility, while psychological states influence how these attributes are evaluated and the threshold against which utility is compared. Experiences with mobility, plug-in requirements, charging and discharging, and other V2G participants can subsequently affect future participation decisions.

Driver preferences are empirically grounded in the stated-choice experiment by Bakhuis et al. (2025). Each driver is assigned to one of four latent preference classes and receives individual sensitivities to financial compensation, guaranteed range, minimum plug-in time, and battery degradation. These




Figure 4.3: Conceptual framework of V2G participation and behavioural feedback. Contract attributes and psychological states determine contract utility and the participation decision; V2G experience in turn feeds back into future decisions.

differences produce heterogeneous evaluations of the same offered V2G contract. The class characteristics and relative importance of the contract attributes are reported in Appendix C.4.

Three psychological states can change through experience. Range anxiety is updated while a vehicle is driving and responds to its resulting state of charge, with low-battery and empty-driving events producing stronger increases. Plug-in time concern is updated once per simulated day by comparing the accumulated plug-in duration with the contractual requirement or a baseline duration for non-participants. Environmental self-identity is evaluated after charging and discharging and changes only when an energy exchange has actually occurred, depending on whether that action is classified as grid-supportive. The model also contains a peer-influence mechanism based on the prevalence and apparent satisfaction of participating peers; this mechanism was disabled in the final analysed experiments.

Changes in the psychological states are applied gradually through an inertia mechanism that retains 90% of the previous value and incorporates 10% of the experience-based adjustment. Once per simulated day, the current psychological states are incorporated into the driver-specific contract preferences and participation threshold. Once per simulated week, drivers re-evaluate the offered contract and participate when its utility exceeds their effective participation threshold. Participation is therefore endogenous and can evolve in response to accumulated mobility, charging, and V2G experiences. The detailed update rules are provided in Appendix C.4.

4.7. Energy balance and grid indicators

The model translates charging and discharging behaviour into aggregate electricity-system indicators. At each tick, household demand, service-sector demand, solar PV production, and wind production are updated from external time-series profiles. Together, these variables define the background electricity-system conditions before EV charging and V2G discharging are added.

The baseline energy balance is calculated as renewable supply minus non-EV demand:

$$\text{baseline-energy-balance}_t = PV_t + Wind_t - Demand_{residential,t} - Demand_{service,t}$$

A positive value indicates that renewable production exceeds building demand at that tick, while a neg-

ative value indicates that building demand is higher than renewable production. This balance is used as an internal signal for flexible charging decisions, because it indicates whether additional charging can be accommodated from the perspective of renewable availability and background demand.

For congestion evaluation, the model uses current net load. This indicator includes building demand, EV charging demand, V2G discharging supply, and renewable production:

$$\text{current-net-load}_t = \text{Demand}_{EV,charge,t} - \text{Supply}_{EV,discharge,t} + \text{Demand}_{residential,t} + \text{Demand}_{service,t} - PV_t - Wind_t$$

A higher current net load indicates greater pressure on the local electricity system. EV charging increases this value, while V2G discharging and renewable production reduce it. This indicator is therefore used both as a system output and as a control signal for peak-shaving behaviour.

The main grid indicator in this study is peak-load exceedance. A congestion threshold is used to identify moments when current net load becomes too high. These exceedances are later translated into the congestion-related loss function described in the methodology chapter. In addition to the loss value, the analysis considers related indicators such as the timing and magnitude of net-load peaks, EV charging demand, V2G discharging supply, and the share of time steps above the congestion threshold.

4.8. Data sources and inputs

The model uses several external data sources and input parameters to represent travel behaviour, vehicle characteristics, charging infrastructure, electricity demand, renewable production, and V2G contract design. Most baseline data inputs are retained from the original model structure.

Travel behaviour is based on daily mobility profiles that determine when EV drivers travel and which type of destination they visit (Centraal Bureau voor de Statistiek, 2016). These profiles are used to assign each driver a sequence of trips for each simulated day. Vehicle characteristics are based either on a fixed battery-capacity setting or on an imported EV fleet composition. In the latter case, vehicle models and battery capacities are sampled from an external updated fleet dataset, allowing heterogeneity in battery size across agents.

Electricity demand and renewable production are represented using external time-series profiles from one specific year (Claessen et al., 2014; Deru et al., 2011; Elia, 2018; KNMI, 2018; Voulis et al., 2017). Household electricity demand, service-sector electricity demand, solar PV production, and wind production are updated at each model tick. These profiles provide the baseline demand and production patterns before EV charging and discharging are added. Household demand is scaled by the number of participating households, while renewable production is scaled by the renewable coverage parameter and the solar-versus-wind ratio. Together, these parameters determine the size and composition of local renewable generation in the model.

The V2G extension adds contract and behavioural input parameters. Each simulation run uses one V2G contract configuration, defined by financial compensation, guaranteed range, minimum plug-in time, and battery degradation. These attributes are used in the utility-based participation decision. Driver preference heterogeneity is represented through latent preference classes based on the stated-choice results used in the Bakhuis et al. extension. Additional psychological parameters define the initial distributions and update rates for range anxiety, plug-in concern, and environmental self-identity.

4.9. Model assumptions and simplifications

Several assumptions and simplifications were made to keep the model focused on local congestion mitigation, endogenous V2G participation, charging-strategy comparison, and system-design exploration. First, each simulation run includes one V2G contract configuration. Drivers therefore decide whether to accept or reject the offered contract, rather than choosing between multiple competing contracts within the same run. This allows contract attributes to be treated as design parameters in the inverse-modelling setup and makes it possible to analyse how specific combinations of financial compensation, guaranteed range, minimum plug-in time, and battery degradation compensation affect participation and congestion outcomes.

Second, the aggregator is assumed to be active whenever V2G is enabled. Connected V2G vehicles transfer charging and discharging control to the aggregator, within the limits imposed by state of charge, charger availability, mobility needs, and guaranteed range. The aggregator uses rule-based control based on residual energy balance and current net load, rather than formal optimisation, perfect foresight, or day-ahead scheduling. The model therefore represents heuristic peak-shaving and grid-supportive charging behaviour, not optimal market bidding or predictive dispatch.

Third, market mechanisms are simplified. Aggregator profit, wholesale prices, retail tariffs, and imbalance settlement are not modelled explicitly. Charging and discharging are instead guided by local system indicators such as residual energy balance and current net load. This reflects the focus of the study on physical system outcomes, especially peak-load reduction, charging flexibility, and V2G availability, rather than on the business case for aggregation.

Fourth, the electricity system is represented at an aggregate local level. Residential demand, service-sector demand, solar PV production, wind production, EV charging demand, and V2G discharging supply are included, but detailed grid topology, voltage constraints, transformer loading, and power-flow effects are outside the model scope. Peak-load outcomes should therefore be interpreted as aggregate net-load indicators rather than location-specific distribution-grid congestion results.

Fifth, non-V2G charging strategies are represented through simplified rule-based policies. Net-aware charging follows a fixed time-of-day power limitation, while smart charging uses rule-based responses to mobility needs and grid conditions rather than optimisation of individual charging sessions. These simplifications keep the comparison between normal charging, net-aware charging, smart charging, and V2G transparent and computationally feasible.

Sixth, battery degradation is included as a contract attribute affecting participation, but not as a physical ageing process. It influences how attractive a V2G contract is to drivers, but does not reduce battery capacity, alter charging efficiency, or directly constrain cycling during the simulation. Energy conversion losses associated with EV charging and discharging are not explicitly represented. Consequently, energy flows between the grid and vehicle batteries are treated without efficiency losses, which may slightly overestimate usable V2G energy and charging efficiency.

Finally, driver and charging-infrastructure behaviour are simplified. Drivers follow daily travel profiles and charging decisions are based on location, charger access, state of charge, charging policy, and participation status. Charger availability is represented at the patch level, but detailed queuing behaviour, location-specific charging prices, connector compatibility, strategic route choice, and strategic driver behaviour are not included. Behavioural adaptation is limited to the implemented feedback mechanisms: range anxiety, plug-in concern, environmental self-identity, and peer influence. Runs with excessive empty-driving behaviour are treated as implausible during the analysis rather than as desired model outcomes.

These simplifications mean that the model should be interpreted as a behavioural and aggregate system-level simulation of V2G participation, charging flexibility, and peak-load mitigation, rather than as a detailed operational model of electricity markets, user route choice, battery ageing, or distribution-grid power flows.

5

Results

5.1. Model verification and validation results

This section summarises the outcomes of the procedures described in Section 3.4. Detailed diagnostics and validation comparisons are reported in Appendix D. Results concerning loss-function selection and stochastic stability are presented separately.

5.1.1. Verification results

The verification procedures did not identify unresolved implementation problems in the final model version. Energy accounting, battery and power constraints, charging-location logic, guaranteed-range protection, and the prevention of simultaneous charging and discharging operated as intended. The Python–NetLogo pipeline also preserved the parameter settings, random seeds, reporters, and calculated losses required to trace each simulation. Detailed criteria and outcomes are reported in Table D.1.

5.1.2. Validation results

The model produced interpretable daily and seasonal patterns in mobility, charging, renewable generation, and net load. Residential charging was concentrated after evening arrivals, workplace charging around morning arrivals, and net-aware charging reduced demand during the constrained evening period. These patterns were broadly consistent with the reference profiles reported by ElaadNL (Refa et al., 2023).

Mobility outcomes indicated that the charging controls generally preserved vehicle availability: empty-driving accounted for less than 0.239% of kilometres in 95% of validation runs. Smart-charging and V2G profiles were assessed qualitatively because their form depends strongly on the implemented control strategy. Their temporal responses were broadly consistent with the intended congestion-oriented behaviour and with the scenarios reported by ElaadNL (ElaadNL, Team Marktonontwikkeling, 2025). An expert review with Jolijn van Dijk of EVConsult further supported their use for comparative analysis.

Overall, the model was considered sufficiently credible for comparing V2G participation, charging strategies, and congestion-related outcomes, but not for predicting the operation of a specific future distribution grid.

5.2. Exploratory results

The experimental analysis developed iteratively through repeated cycles of simulation, inspection, and refinement. The 19 experiment series, comprising 4,510 individual runs and approximately 890 unique parameter combinations, were therefore not all designed for the same purpose. Early series broadly explored model behaviour and alternative parameter spaces, while other series focused on V2G control inputs, reduced parameter spaces, or corrections to the Bayesian sampling procedure. Later series used larger numbers of repetitions to examine stochastic stability. Most parameter combinations were repeated three or five times, whereas dedicated stability experiments used 50 repetitions. An overview

of the experiment series, their purpose, and their repetition structure is provided in Appendix Table E.1.

The exploratory results are structured around the main decisions required for the final experiments and analysis. First, alternative performance measures are compared to explain the selection of the 600 kW squared loss function. The development of the Bayesian sampler across experiment series is then assessed, followed by an evaluation of the number of repetitions required for stable results. Finally, the model outputs and parameter effects are examined to determine which outcomes and parameters should be retained for the subsequent analysis and whether the selected parameter space produces sufficiently stable and interpretable results. Where necessary, experiment series that were not directly comparable because of differences in outputs, loss definitions, or preprocessing were excluded from specific comparisons.

5.2.1. Run Example

Before introducing the exploratory results, Figure 5.1 illustrates how the different electricity-flow components combine into the local net-load profile during an example simulation day.




Figure 5.1: Daily profile of electricity flows - day 22

Building demand remains relatively high throughout the day and peaks in the evening, while renewable generation is more irregular and concentrated around midday. Grid-aware charging also increases during the morning and evening. From around 17:00, charging power is reduced, visible as a sharp drop, although continued vehicle arrivals keep aggregate charging demand elevated. V2G responds to the residual system balance: vehicles charge when sufficient headroom is available, particularly during lower-demand periods in the morning or periods of high renewable generation around midday, and discharge during the evening peak to reduce net load.

5.2.2. Loss function

Alternative loss formulations were compared post hoc to determine which specification remained informative across both poorly and well-performing regions of the parameter space. The measure also needed to capture the additional pressure created by EV charging and V2G operation during already elevated demand, rather than primarily reflecting the underlying household and service-demand profile. The comparison focused on four properties: the ability to distinguish between parameter combinations, consistency in the relative ranking of runs, sensitivity to extreme exceedances, and stability across stochastic repetitions. The results below explain why the 600 kW squared-exceedance loss was selected for the subsequent analyses. A complete overview of the tested formulations and supporting diagnostics is provided in the Appendix Section E.2.

Figure 5.2 illustrates the two main components of the selected loss formulation. Panel a places the

selected 600 kW threshold in relation to the average daily net-load profiles, while Panel b shows how the squared penalty gives progressively greater weight to larger exceedances than a linear penalty, without concentrating as strongly on only the most extreme events as the higher-order alternatives.




Figure 5.2: Illustration of the selected congestion-loss formulation. Panel (a) shows average daily net-load profiles by month across exploratory runs with renewable coverage 40%, together with alternative congestion thresholds θ . Panel (b) compares the normalised penalty assigned to exceedances above the selected 600 kW threshold for different penalty exponents p .

Initially, a threshold of 800 kW appeared suitable, because it focused attention on the highest net-load periods. However, as the Bayesian sampler started identifying better-performing regions of the parameter space, this threshold became less informative. As illustrated by the daily profiles Figure 5.2a shows even the highest monthly average peak loads remained well below 800 kW. Many improved runs therefore no longer exceeded the 800 kW threshold, or only exceeded it very rarely. As a result, many runs received a loss of zero or near zero, making it difficult to distinguish between parameter combinations that were all “good enough” under this threshold. The descriptive effect of the threshold on zero losses and the loss distribution is reported in Table E.2.

The rank-correlation analysis showed that alternative penalty forms at the same threshold generally ranked parameter configurations very similarly. Substantial agreement was also retained between the 500 and 600 kW formulations, indicating that these thresholds broadly identified the same regions as relatively favourable or unfavourable. The 700 kW formulations were likewise highly correlated with one another, but showed weaker agreement with the lower-threshold formulations. This indicates that the choice of threshold affected configuration ranking more strongly than the precise penalty order within a given threshold. Descriptive relationships with the sampled parameters showed the same general pattern, while becoming less informative at higher thresholds as increasing numbers of runs received zero or near-zero loss. A representative correlation matrix is shown in Figure E.1.

Higher-order penalties, particularly powers of 10 and 50, increasingly concentrated the loss on a small number of extreme exceedances. As illustrated conceptually in Figure 5.2b, the penalty becomes increasingly concentrated near the largest exceedances as the power increases. This produced highly skewed numerical distributions, with many values close to zero and a small number of disproportionately large outcomes. Although these formulations sometimes appeared relatively stable, this was partly caused by the limited variation among the many near-zero runs rather than by consistently robust performance. They were therefore less suitable for distinguishing between parameter combinations across a broader range of peak conditions. Similarly, thresholds of 650 kW and above frequently produced zero or near-zero values, which limited their usefulness for comparison.

Stochastic stability differed between the loss formulations, but did not identify one uniquely superior specification. Across the practically relevant thresholds between 400 and 600 kW, the count-based, linear, squared, and cubic formulations showed variability of broadly the same order of magnitude, with similar patterns across experiment series. Relative variability generally increased at higher thresholds as the loss distributions became more concentrated near zero and increasingly right-skewed. The 600 kW squared loss therefore showed somewhat greater dispersion than lower-threshold alternatives, but its stochastic variability remained sufficiently small relative to the differences between parameter




Figure 5.3: Median proposal loss across consecutive Bayesian experiment series during the exploratory phase. Each point is the median loss of one proposed parameter combination across its repetitions. Series labelled Untrained started without a fitted surrogate, and the vertical line marks the correction of the sampling procedure. The y-axis is logarithmic.

configurations. By contrast, the higher-order penalties showed substantially greater relative variability, consistent with their stronger dependence on a small number of extreme exceedances. The comparison across thresholds and penalty forms is shown in Figure E.2.

Overall, the 600 kW squared-exceedance loss was selected as a practical compromise rather than as a uniquely optimal formulation. The average building-demand profile already approaches 600 kW during its recurring peak, so exceedances above this level were interpreted as additional pressure on the upper system peak, rather than as ordinary electricity use. At the same time, the threshold retained sufficient non-zero variation to distinguish between parameter combinations. The squared penalty gave greater weight to more severe exceedances while remaining sensitive to a broader range of peak events, without allowing a small number of extremes to dominate the overall result. The selected formulation therefore matched the objective of identifying configurations that create meaningful headroom during periods of elevated net load and remained applicable across the different charging and V2G control strategies.

5.2.3. Learning behaviour of the sampler

The exploratory experiments were generated using a Bayesian sampling procedure rather than as one static random sample of the parameter space. The sampler used previously observed losses to select new parameter combinations, meaning that the order of the runs is relevant for interpretation. The learning behaviour of the sampler is therefore itself an important exploratory result. If the sampler does not improve over time, the results would mainly reflect random search. If it does improve, later batches can be interpreted as increasingly focused on regions of the parameter space associated with lower congestion-related loss.

Figure 5.3 shows the development of proposal-level median loss across the exploratory Bayesian experiment series. Appendix Table E.3 provides complementary experiment-series statistics based on the individual runs. Because the series differed in parameter space, sampler configuration, training history, and experimental purpose, they do not represent a single uninterrupted learning trajectory. Instead, the figure shows how proposal quality developed within and across successive stages of the exploratory sampling process. Within each experiment series, four parameter combinations were evaluated before the sampler was updated, so learning occurred batch-wise rather than after every individual proposal.

The trajectory shows a clear difference between the early diagnostic series and the later corrected Bayesian runs. In the early series, the median losses were generally higher and more variable, indicating that the sampler had not yet consistently identified useful regions of the parameter space. Also in the table we clearly see larger differences between the median, upper quartile, maximum, and MAD, indicating that they contained both low-loss and very high-loss individual runs. Minimum losses could nevertheless already be similar to those observed later, showing that favourable configurations were

occasionally found early; an isolated low-loss run, however, does not indicate that the sampler could identify such configurations consistently or that the underlying parameter combination performed reliably across repetitions. An issue in the earlier configuration in these earlier runs was that the sampler configuration attempted to combine low loss with high variation. This was unintended and not useful, because high-loss outcomes also tended to be more variable. After the sampler was corrected and fixed, the trajectory shifted towards lower values and the later runs show a stronger concentration of median losses in the lower-loss region.

Some earlier experiments used the dynamic sampling setting, whereas later experiments increasingly relied on exploitation. The exploration–exploitation balance was not systematically tuned, so these results should not be interpreted as a formal comparison between sampler strategies. However, lower-loss parameter combinations also tended to show lower repetition-level variability, indicating no clear trade-off between congestion performance and stochastic stability. This relationship is examined in more detail in Section 5.2.4. It also reduced the concern that stronger exploitation would favour less stable configurations. Given the objective of identifying low-loss regions within a limited budget of computationally expensive full-year simulations, the exploitative setting therefore appeared more suitable and was retained for the final learning phase.

Repeated selection of exactly the same parameter combination depended on the size of the parameter space, the simulation horizon, and the sampler setting. In the full-year runs with a broad parameter space, repeated combinations did not occur because the number of possible combinations was large. In shorter and more restricted experiments, repeated combinations occurred more often because fewer parameters were varied and the simulated period was less variable. For example, when only a winter period was simulated, the weather conditions and system stress were more consistent than across a full year. This made it easier for the sampler to repeatedly identify similar or identical promising settings. This was useful diagnostically, because it showed that the sampler could concentrate on parameter regions that performed well under specific conditions.

Overall, the sampler became more effective as the experimental setup and sampling procedure were refined. Early series primarily served diagnostic and exploratory purposes, whereas the later corrected and more exploitative series produced lower and more concentrated median losses and were therefore more informative for identifying promising parameter regions.

5.2.4. Repetitions, stochasticity and stability

The stability analysis was used to assess how many repetitions are needed to obtain a reliable estimate of model performance. Because the model contains stochastic elements, the loss of a parameter combination can differ between repetitions. This is especially relevant for a threshold-based loss function: small changes in simulated behaviour can determine whether net load remains just below the threshold or slightly exceeds it, which can create relatively large differences in loss. The stability analysis therefore examined whether conclusions based on a limited number of repetitions remain consistent when compared with a 50-repetition reference.

The analysis was based on six experiment series with 50 repetitions: four later stability runs and two earlier stability-oriented runs. Subsamples of 5, 10, 20, 30, and 40 repetitions were compared with the full 50-repetition estimate. Figure 5.4 shows the Spearman rank correlation between each subsample estimate and the 50-repetition reference. Spearman rank correlation was used because the main question is whether parameter combinations are ranked in a similar order, rather than whether the exact loss values are identical.

The results show that central performance stabilises quickly. For median loss, the rank correlation with the 50-repetition estimate is already high with five repetitions and increases further with ten repetitions. This suggests that a limited number of repetitions is sufficient to distinguish broadly better and worse parameter combinations in terms of central loss performance. The same pattern is visible for the combined measure, although this measure is partly influenced by variability.

MAD is less stable than median loss. Its rank correlation with the 50-repetition estimate is much lower at five and ten repetitions, and only becomes clearly stronger after more repetitions. This is expected because MAD measures variability itself, and variability requires more observations to estimate reliably than central performance. MAD-based conclusions are therefore useful as supporting evidence, but




Figure 5.4: Spearman rank correlation between performance estimates based on subsamples of repetitions and the corresponding 50-repetition reference. The horizontal axis shows the number of repetitions used (5–50); separate lines show median loss, MAD, and the combined measure.




Figure 5.5: Relationship between central performance and stochastic variability across parameter combinations. Each point is one parameter combination; the horizontal axis shows median loss and the vertical axis the MAD across repetitions (both on a logarithmic scale).

should be interpreted more cautiously when the number of repetitions is limited.

The distributions of loss across repetitions were often skewed, with a small number of high-loss repetitions pulling the mean upward relative to the median. This illustrates why median loss provides a more robust measure of central performance, while MAD was used to describe variation across repetitions.

The stability analysis also showed that stochastic variability was not evenly distributed across the parameter space. Parameter combinations with higher median loss generally also showed higher variation across repetitions, as shown in Figure 5.5. This means that poorly performing configurations were often not only worse centrally, but also more variable across repetitions.

Part of this relationship follows from the formulation of the power-based loss: larger exceedances increase both the typical loss and the potential absolute variation between repetitions. This pattern was less apparent for the exceedance-count loss, which records only whether the threshold is exceeded and not the magnitude of the exceedance. As shown in Appendix Figure E.3, parameter combinations with similar median exceedance counts could therefore show substantially different relative variability.

Importantly, the results did not indicate a trade-off in which low median loss could only be achieved at the cost of high variability. Low-loss configurations generally also showed relatively limited MAD. A single run is therefore especially unreliable for parameter combinations in the higher-loss region.

Overall, ten repetitions provided a practical compromise for the subsequent analysis. Median loss was sufficiently stable to distinguish broadly better and worse parameter combinations, while keeping the computational burden manageable. Ten repetitions were less suitable for strong claims about robustness, so MAD is interpreted more cautiously. Lower-loss configurations also tended to have lower MAD, meaning that the desired outcomes were generally also more stable, although this pattern is partly driven by the squared loss formulation itself.

5.2.5. Outcome selection and exploratory output relationships

Before interpreting parameter effects, the exploratory outputs were compared to determine which outcomes were most useful to analyse alongside loss. This step was necessary because several variables describe related aspects of model behaviour and are therefore not independent. Analysing all outputs separately would add unnecessary repetition and could give too much weight to closely related indicators. The aim was therefore to identify a smaller set of relevant output groups that capture different dimensions of the model.

Correlations were calculated between the aggregated output measures, including median, MAD, and combined values. In addition, the median, mean, and final value of each output were compared with one another to assess whether important temporal dynamics might be obscured by aggregation and




Figure 5.6: Exploratory Spearman rank correlations between loss and selected aggregated model outputs. The matrix was used to identify redundant indicators and select output groups for further analysis.

therefore required separate inspection. A first pattern was that the combined measure was generally very strongly correlated with the corresponding median. This indicates that, for most outputs, the combined measure was mainly driven by central performance rather than providing a fully separate outcome. MAD showed a different pattern: for most outputs, it was only weakly correlated with the corresponding median and combined measures. Loss was an exception, where MAD was more strongly related to central performance. This is consistent with the earlier observation that higher-loss configurations also tended to be more variable.

Figure 5.6 presents the correlations between loss and the other aggregated model outputs. For each parameter combination, the selected outputs were aggregated across repetitions using the median. V2G participation and SOC were first averaged over the simulation horizon, while charging and discharging outputs were summed over the full run. The correlations with loss showed that V2G participation was not a sufficient explanation for congestion performance. V2G-count was only weakly correlated with loss, meaning that higher participation did not automatically imply lower congestion-related loss. EV charging demand showed a moderate positive association with loss, suggesting that charging behaviour contributes to congestion pressure but does not fully determine it. This supports the interpretation that the timing and availability of flexibility matter more than participation count alone.

Several charging-related outputs were strongly correlated with each other. V2G-count was strongly positively correlated with V2G charging demand and strongly negatively correlated with non-V2G charging demand. This mainly reflects a shift in charging from non-V2G vehicles to V2G vehicles as more vehicles participate. V2G discharging was also related to charging demand, but these relationships should be interpreted as indicators of coupled model behaviour rather than as direct performance effects.

One of the strongest exploratory relationships with loss was found for battery-state variables. Loss was strongly negatively correlated with the SOC difference between V2G and non-V2G vehicles, and also with the absolute SOC of V2G vehicles. In other words, runs in which V2G vehicles had lower SOC values tended to have higher loss. This suggests that congestion performance was more closely related to the battery state and available discharge flexibility of participating vehicles than to participation count

alone.

Finally, the exploratory parameter-relevance methods were interpreted cautiously because their results depended on the selected experiment series, preprocessing choices, and relationships between the outputs themselves. The permutation-importance results showed that strongly correlated outputs were often influenced by similar input parameters, suggesting that some observed output relationships may reflect shared underlying drivers. SHAP patterns were particularly sensitive to which experiment series were included and to whether high-loss runs were filtered out. Statistical differences between parameter groups also became clearer as more trained data were added, although not every input variable showed significant or consistent differences. These analyses were therefore used mainly as screening tools to identify relevant outputs, parameters, and relationships for the subsequent analysis, rather than as definitive evidence of parameter effects.

Overall, the exploratory output comparison showed that several model outputs contained strongly overlapping information. V2G participation alone was only weakly related to loss, while charging and discharging measures mainly captured related aspects of EV behaviour. Battery-state variables provided additional information, with lower V2G SOC generally associated with higher loss. The subsequent analysis therefore focuses on a smaller set of complementary outcomes representing participation, charging behaviour, and available V2G flexibility.

5.2.6. Contract and control settings selected for further analysis

The exploratory runs were also used to refine the contract settings for the subsequent analysis. In the early experiments, contract parameters were allowed to vary over ranges that extended well beyond the values included in the stated choice experiment on which the latent classes were based. Participation estimates for these extreme combinations were therefore less reliable and often pushed participation towards its maximum, reducing sensitivity to differences between individual contract attributes. For the later analysis, the contract ranges were restricted to values closer to the original choice experiment, while still covering almost the full participation range, from 0 to 495 V2G participants. This reduced the influence of poorly supported extreme contracts and improved the interpretability of the resulting parameter effects.

The V2G control thresholds were also refined during the exploratory phase. Several threshold combinations were examined through repeated runs and visual inspection of the resulting load profiles. These checks were combined with tests using different loss thresholds. The aim was to identify settings that reduced congestion-related loss consistently across repetitions, while avoiding excessive discharge and limiting the distance driven with an empty or very low battery.

The aggregated threshold comparisons in Appendix Table E.4 show that congestion loss, stochastic variability, and mobility outcomes did not improve in parallel. Threshold values associated with lower loss sometimes produced substantially more driven-empty outcomes, while settings with more favourable mobility outcomes were not always those with the lowest loss. The selection therefore also considered which combinations of control thresholds performed well across runs, together with SOC behaviour and V2G participation. Since each threshold value was evaluated in combination with different settings of the other control parameters, the aggregated comparisons were interpreted as indicative rather than as isolated effects of individual thresholds.

Based on these exploratory runs, the selected V2G control configuration was `charge-reduce-start-kW = 450`, `charge-stop-kW = 650`, `discharge-start-kW = 550`, and `full-discharge-kW = 750`. The configuration was selected as a practical compromise rather than as a fully optimised control rule. It reduced charging before it intensified existing peaks, while retaining sufficient opportunity for vehicles to charge. Discharging was activated before the highest net-load periods, but full discharge was reserved for more severe conditions to reduce the risk of depleting batteries too early. This was particularly important because the rule-based controller did not anticipate future peaks.

Overall, the exploratory phase narrowed the contract ranges and established a workable V2G control configuration for the final experiments. The revised contract space remained broad enough to capture large differences in participation, while avoiding poorly supported extreme combinations. The selected control thresholds balanced congestion reduction, stability across repetitions, and mobility effects, and should therefore be interpreted as fit-for-purpose settings rather than an optimised control strategy.

5.3. Final results

This section presents the final results from the two independently executed learning chains. It first examines how the sampler developed across and within learning rounds and whether performance improved and stabilised over time. The analysis then considers the temporal composition of the remaining congestion loss, identifies the parameters and parameter regions associated with low-loss configurations, and evaluates how these configurations relate to other relevant system outcomes. A complete overview of the experiment series and realised simulation runs is provided in Table F.1.

5.3.1. Learning behaviour of the sampler

Sampler learning was assessed both between successive learning rounds and within experiment series across internal Bayesian batches. The analysis compares the two learning chains and their paired parallel experiment series to determine whether proposal quality improved and whether this improvement was consistent across search trajectories.

Learning between rounds

Sampler performance was evaluated across seven learning rounds in both learning chains. As described in methodology Section 3.5, between-round performance was assessed using the pooled distribution of proposal-level mean losses from the two parallel experiment series.

Figure 5.7 and Table F.2 show that the largest improvement occurred after the first learning round. Median proposal loss decreased from 2.835 to 0.978 in V2G 1 and from 2.417 to 0.719 in V2G 2 between rounds 1 and 2, corresponding to reductions of 65.5% and 70.3%, respectively. Subsequent changes were smaller and were not monotonic. V2G 1 improved further in round 3, followed by alternating improvement and deterioration in later rounds. V2G 2 followed a more stable trajectory, improving further until round 4 and showing only relatively small changes thereafter.

By round 7, median proposal loss was approximately 77% below its initial value in both learning chains. V2G 1 showed greater differences between its paired parallel experiment series, most clearly in round 4, when their median losses were approximately 0.5 and 1.5. V2G 2 showed smaller differences between its parallel series after the first two rounds and remained more consistently concentrated at lower loss levels. The overall improvement was therefore substantial in both learning chains, but most of it occurred during the early rounds.

The main change in both learning chains was the reduction of poor proposals (Figure 5.8). In V2G 1, the share of proposals with loss above 2.0 decreased from 57.5% to 18.2%, while V2G 2 eventually eliminated this high-loss category entirely. At the same time, the share of low-loss proposals increased in both chains, with V2G 2 showing the stronger concentration in the low-loss region. Overall, the reduction of poor proposals was more pronounced than the improvement among the very best proposals. The sampler therefore learned mainly to avoid unfavourable regions rather than to concentrate all proposals near a single optimum. Round 7 did not produce a further substantial shift in this pattern.

This pattern is also reflected in the best-performing proposals. Low-loss configurations were identified early in both learning chains: the cumulative best mean proposal loss improved substantially by round 2 and changed only slightly thereafter. The performance of the best 20% of proposals showed a similar pattern, indicating that a favourable region had already been reached in the early rounds. V2G 1 reached this plateau somewhat earlier, while V2G 2 continued to make small improvements. This confirms that later learning mainly refined the search within favourable regions and excluded high-loss configurations rather than identifying substantially better best-performing configurations. The remaining gains therefore diminished over time, although this does not establish formal convergence.

Learning across internal Bayesian batches

Learning within the main learning rounds was assessed by comparing the median proposal loss of the first two and final two internal Bayesian batches (Table F.3). Clear reductions occurred in V2G 1 rounds 1, 3, and 5 and in V2G 2 rounds 1, 2, and 4. Other rounds showed limited improvement or deterioration, particularly V2G 1 rounds 4, and 7, and in V2G 2 rounds 6 and 7. Learning within experiment series was therefore less consistent than the improvement observed between successive learning rounds.

The paired parallel experiment series did not always follow the same within-round trajectory, experiment series often showed the same broad direction, but their consistency differed between model versions.




Figure 5.7: Median proposal loss across successive Bayesian learning rounds for (a) V2G 1 and (b) V2G 2. Each point represents one proposed parameter configuration, aggregated across its repetitions; proposals are grouped by learning round (R1–R7) and parallel experiment series (A and B). The legend reports the relative change in proposal-level performance between consecutive learning rounds (B) and within each learning round (W), with positive values indicating improvement through lower loss. Loss is shown on a logarithmic scale.




Figure 5.8: Composition of proposal quality across the main learning rounds for V2G 1 and V2G 2. Each proposal is assigned to one mutually exclusive loss category based on its median loss over ten repetitions.

The development of proposal loss across internal Bayesian batches is shown separately for both learning chains and their paired parallel experiment series in Figure F.1. In round 7, one series improved and one deteriorated within each learning chain. V2G 1 showed larger differences between parallel runs, most clearly in round 4. The V2G 2 runs were generally more similar after the first round, although isolated high-loss proposals still produced temporary deviations. These comparisons are descriptive, as each learning round contained only two parallel runs.

Stopping criterion

The predefined stopping criterion was not fully met by the end of the seventh learning round. Proposal loss had decreased substantially relative to the initial rounds, but the later-round changes remained somewhat unstable. In the final rounds, changes in Q25 and median loss were below 10% for both learning chains, indicating that the central part of the loss distribution had largely stabilised. The Q75 was less consistent: the required threshold was met in round 6 but not any more in round 7. Because all conditions had to be satisfied for two consecutive rounds, the formal stopping criterion was not reached.

Overall, the clearest evidence of learning was observed between successive learning rounds, with the largest improvement occurring after the first update. Within-round learning was also visible, but was less consistent and more sensitive to individual search trajectories. Across both learning chains, favourable low-loss regions were identified early, while later rounds mainly reduced the occurrence of poor proposals and refined the search within these regions. Improvements therefore became smaller and non-monotonic over time. Although the formal stopping criterion was not reached, the later rounds showed substantially more stable central performance and diminishing additional gains.

5.3.2. Temporal composition of congestion loss

Congestion loss was highly concentrated in a limited number of five-minute intervals (Figure 5.9a). Across the analysed configurations, 80% of total loss was generated during 1,871 ticks, equivalent to approximately 156 hours or 26% of all exceedance ticks. The remaining loss accumulated more gradually through a larger number of smaller contributions, particularly among the worst-performing 20% of configurations. Loss was therefore distributed unevenly across exceedance events rather than being generated uniformly whenever the congestion threshold was exceeded. This concentration was even stronger among the best-performing configurations. Although their total loss was substantially lower, a relatively large share of the remaining loss was generated by their largest residual peaks. This suggests that better-performing configurations eliminated many smaller and intermediate exceedances, while a limited number of more difficult peak events remained.

The daily distribution shows that congestion loss occurred predominantly during the evening, with the largest contribution between approximately 17:00 and 21:00 (Figure 5.9b). A much smaller secondary contribution occurred between 09:00 and 11:00 at the best-performing configurations. Both performance groups followed the same broad evening-peak pattern, but the worst-performing configurations produced substantially larger and somewhat more prolonged losses.

Across the simulation horizon, loss was concentrated mainly near the beginning and end of the simulated period, corresponding to the winter-related periods (Figure 5.9c). This seasonal concentration is consistent with the modelled winter conditions, when higher background demand and lower solar generation leave less headroom before the congestion threshold is exceeded. The simulations did not cover a complete calendar year, and the final, incomplete month was excluded from the monthly plot. Congestion was not confined entirely to these periods, particularly among the worst-performing configurations, which continued to generate loss throughout much of the intermediate part of the year. Nevertheless, the best- and worst-performing groups followed a broadly comparable seasonal pattern. Better-performing configurations primarily reduced the magnitude of winter-related exceedances rather than shifting congestion to another part of the simulated year.

A comparison of system designs across seasons showed that relative performance was largely persistent (Figure F.6). Of the configurations in the best-performing winter group, 67% were also among the best-performing configurations in summer, while 85% of the middle group and 81% of the worst-performing group remained in the corresponding performance category. Seasonal conditions therefore changed the ranking of some configurations, but did not fundamentally reverse the overall performance pattern.




Figure 5.9: Temporal composition of congestion-related loss for the best- and worst-performing parameter configurations. Panel (a) shows how loss is concentrated across exceedance ticks, panel (b) shows its distribution over the day, and panel (c) shows its distribution across the simulated horizon. The simulations did not cover a complete calendar year; the final incomplete month was omitted from panel (c).

Overall, the results show that congestion loss was driven mainly by a small number of winter evening peaks rather than by persistent exceedances throughout the simulation. Better-performing configurations reduced the frequency and magnitude of exceedances across most periods, leaving only a limited set of residual peak events. Their relative performance was largely maintained across seasons, indicating that favourable configurations were not specific to one seasonal period.

5.3.3. Parameters associated with congestion loss

To identify which model inputs were most useful for explaining congestion loss, tree-based surrogate models were fitted to the aggregated configuration-level outcomes. Feature scoring was then used to assess how strongly each parameter contributed to the prediction of loss. These scores indicate predictive importance rather than causal effects or the direction of a parameter's influence.

Surrogate models were trained on both raw loss and log-transformed loss. The out-of-bag R^2 measures how well the model predicts observations that were not used to train the individual trees, and therefore provides an internal estimate of predictive performance. The raw-loss distribution was strongly right-skewed, with a limited number of very high-loss configurations. The log transformation compressed this upper tail, reducing the influence of these extreme observations and producing substantially better and more stable predictive performance. The median out-of-bag R^2 increased from approximately 0.58 for raw loss to 0.89 for log-transformed loss, while variation across model fits also decreased. Log-transformed loss was therefore used as the main target for the feature-scoring analysis, with the raw-loss results retained as a robustness check. The corresponding model-performance comparison is provided in Appendix Figure F.5a.

Renewable coverage and home charging policy were the two dominant predictors of log-transformed loss, with median normalised feature importances of approximately 33% and 25%, respectively. Together, they accounted for more than half of the total feature importance. Guaranteed range was the most important contract-related parameter, with an importance of approximately 9%. The remaining contract, infrastructure, renewable-mix, and charging-policy parameters each contributed a substantially smaller share, generally below 5% for the log-transformed target. Several contract-related parameters received somewhat higher relative feature scores when raw loss was used instead. This is consistent with the greater influence of extreme high-loss observations on the raw-loss models, while renewable coverage and home charging policy remained dominant under both transformations. The broader comparison across raw and log loss, alternative models, and alternative targets is provided in Appendix Figure F.3. Figure 5.10 instead compares the feature-importance patterns for median congestion loss and repetition-level variability.

The relative importance of parameters changed over the learning rounds. Renewable coverage was especially important during the earlier rounds, after which high renewable-coverage values became increasingly common among the better-performing proposals. As the sampler concentrated on this narrower part of the parameter space, renewable coverage provided less variation with which to distin-




Figure 5.10: Median normalised Extra Trees feature importance for log-transformed median congestion loss and log-transformed loss variability (MAD). Bars show the median importance across bootstrap fits, while the horizontal error bars represent the 10th–90th percentile interval.

guish subsequent proposals. This concentration around similar parameter values among the lowest-loss proposals is also visible in Appendix Figure F.4. Home charging policy subsequently explained a larger share of the remaining loss differences. Changes in feature importance over learning rounds should therefore not be interpreted as changes in the underlying physical importance of a mechanism alone, because importance also depends on how much variation in each parameter remains in the sampled data. The development of feature importance across learning rounds is shown in Appendix Figure F.2.

A separate analysis examined the median absolute deviation of loss across repetitions. The Extra Trees models explained this repetition-level variability reasonably well. Median out-of-bag R^2 values were approximately 0.83 for raw MAD and 0.82 for log-transformed MAD. Although their median performance was similar, the raw-MAD models were considerably less stable across fits (Appendix Figure F.5b). The same main parameters that predicted median loss also explained part of the variation between seeds. Renewable coverage remained dominant, while home charging policy was relatively more important for median performance and minimum plug-in time and private charger availability were relatively more important for variability (Figure 5.10). Nevertheless, the number of repetitions per configuration was selected primarily to obtain a stable median ranking and remained relatively limited for precise estimation of MAD. The variability results should therefore be interpreted as indicative.

Several robustness checks were used to assess whether the main ranking depended on a specific model or outcome transformation. Extra Trees results were compared with Random Forest feature and Univariate scores and produced very similar results especially for log-loss. Models were also repeated using alternative surrogate-model random seeds. The parameter rankings for raw and log-transformed loss were strongly correlated (Spearman's $\rho = 0.742$, $p = 0.0037$) and complete agreement for the three highest-ranked parameters: renewable coverage, home charging policy, and guaranteed range. Differences occurred mainly among the lower-ranked parameters. Overall, the analyses consistently identify renewable coverage, home charging policy, and guaranteed range as the most important predictors of congestion-loss variation, while the following section examines the parameter ranges associated with low-loss outcomes. The complete surrogate-model robustness analysis, including raw-versus-log target performance, alternative tree-based models, univariate scoring, repeated model fits, and learning-round-specific feature scores, is reported in Appendix Figure F.3.

Overall, the feature-scoring analyses consistently identified renewable coverage and home charging policy as the dominant predictors of congestion loss, with guaranteed range as the most prominent contract-related parameter. The ranking was broadly robust across alternative models and target definitions, although the importance of lower-ranked parameters varied more strongly. Several contract-related parameters received relatively greater importance for raw loss than for log-transformed loss, suggesting that they were more informative for distinguishing higher-loss outcomes.

5.3.4. Regions associated with low congestion loss

The feature-scoring analysis identified which parameters contributed most strongly to predicting congestion loss, but did not show which combinations of parameter values were associated with favourable outcomes. Patient Rule Induction Method (PRIM) was therefore used to identify regions of the parameter space in which low-loss configurations were concentrated. In PRIM, *coverage* is the share of target cases captured by a box, while *density* is the share of observations inside that box that belong to the target group. *Mass* represents the share of the complete dataset contained in the box, and *lift* compares the density inside the box with the target-group frequency in the full dataset. A useful box should therefore capture a substantial share of the target cases while remaining sufficiently selective and interpretable. PRIM starts with the full dataset and iteratively removes observations from the edges of the parameter space to create a smaller box with a higher concentration of low-loss cases. It then compares alternative boxes along this peeling trajectory, allowing a trade-off between capturing many target cases and isolating a more selective high-performing region. For parameters explored at discrete values, the resulting box boundaries can fall between tested settings. These boundaries should therefore be interpreted as separating the tested values retained inside and outside the box, rather than as newly evaluated parameter values.

The best-performing 20% was used as the principal target because it represented a clearly low-loss subset while still yielding a well-defined and sufficiently populated PRIM box. For the best 20%, box 9

provided the most balanced combination of coverage and density (Appendix Figure F.10b). It captured 75.6% of the best-performing cases and had a density of 77.5%, while containing 19.6% of all evaluated configurations. Its lift of 3.86 indicates that configurations inside the box were almost four times as likely to belong to the best 20% as configurations drawn from the full dataset. The parameter restrictions of the selected box are visualised in Figure 5.11.

It was characterised by smart home charging, high renewable coverage, a medium-to-high guaranteed range, and a relatively solar-oriented renewable mix. The box also retained moderate-to-high private-charger availability and excluded the highest tested values for minimum plug-in time and battery degradation. The observed loss of configurations inside and outside this region is compared in Figure 5.12.




Figure 5.11: Parameter restrictions of the selected PRIM box for the best-performing 20% of configurations. Grey lines indicate the full explored range and blue lines the range retained within the selected box; only smart home charging was retained for the categorical charging-policy parameter.




Figure 5.12: Observed congestion loss inside and outside the selected PRIM region for the best-performing 20% target. Loss is log-transformed.

Observed outcomes inside the selected box were clearly better than those outside it. Median loss was approximately 0.37 inside the box, compared with 0.90 outside it. The corresponding mean values differed even more strongly, although these were affected by the highly skewed loss distribution. The box should therefore be interpreted as a region with a substantially higher concentration of low-loss outcomes rather than as a deterministic rule: not every configuration inside the box performed well, and some low-loss configurations remained outside it.

The result was robust to several PRIM settings. Varying the peeling parameter and minimum-mass threshold repeatedly selected the same box, with unchanged coverage, density, and restricted dimensions. Resampling also reproduced the principal conditions consistently, the stability of these restrictions under resampling is shown in Appendix Figure F.8. Renewable coverage, solar–wind mix, home charging policy, guaranteed range, and minimum plug-in time were fully recovered in all resamples. The battery-degradation restriction was reproduced in 95% of the samples, while the private-charger restriction was recovered in 76%. This indicates that the first group of conditions formed the most stable core of the low-loss region, whereas private-charger availability was less consistently required.

The box also transferred reasonably well between the two independent learning streams. A box trained on V2G 1 achieved a density of 0.65 and coverage of 0.75 when applied to V2G 2. Conversely, the box trained on V2G 2 achieved a density of 0.84 and coverage of 0.63 when evaluated on V2G 1. Although performance was lower than when boxes were evaluated on their own training stream, both cross-stream applications retained substantial lift and identified meaningful concentrations of low-loss cases. This suggests that the main low-loss conditions were not specific to one learning trajectory. The corresponding cross-chain coverage and density values are shown in Appendix Figure F.9.

The neighbouring PRIM boxes help distinguish the core conditions from these refinements. Box 7 did not restrict battery degradation, while box 8 still allowed the full range of private-charger availability. Both achieved higher coverage but lower density than box 9. Box 10 increased density further by also excluding the 7-hour minimum plug-in option, but this reduced coverage.

The same broad pattern was visible under alternative definitions of favourable performance. The best

10% was less clearly separated, with both lower coverage and lower density despite a similarly restrictive box (Appendix Figure F.10). The very lowest-loss configurations therefore appeared more heterogeneous and were less easily captured within a single compact PRIM region. The best 30% produced a simpler region defined by smart home charging, renewable coverage of 60%, guaranteed range of 105 or 135 km, a solar–wind setting of 50 or 60, and minimum plug-in time below 9 hours. For the best 50%, only three restrictions remained necessary: smart home charging, renewable coverage of 60%, and a guaranteed range of 75–135 km. This progression suggests that smart home charging and high renewable coverage formed the broadest conditions associated with better performance, while guaranteed range and the other restrictions became more selective when the definition of low loss was tightened.

The analysis of poorly performing configurations produced a simpler result. The worst 20% was primarily characterised by home charging that was either normal or net-aware rather than smart. This single restriction captured 73.2% of the worst-performing cases and produced a density of 75.6%, corresponding to a lift of 3.76. Median loss inside this box was approximately 3.95, compared with 0.55 outside it. In contrast, no informative box could be identified for the worst 50%, indicating that the broader lower half of the performance distribution was too heterogeneous to describe using one compact parameter region.

Overall, the PRIM analysis identifies a robust low-loss core consisting of smart home charging, high renewable coverage, and a medium-to-high guaranteed range. A relatively solar-oriented renewable mix, avoidance of the highest minimum plug-in requirement and battery-degradation value, and sufficient private charging availability further refined this region. These conditions should be interpreted as empirically associated design regions within the explored parameter space, rather than as necessary or sufficient causal conditions or as evidence of one unique optimum.

5.3.5. Participation and charging-related outcomes

Having identified the parameter regions associated with low congestion loss, this subsection examines model outcomes beyond the loss metric itself. It first evaluates the relationship between V2G participation and congestion loss, including differences between learning chains and seasonal conditions. Participation is summarised as the mean over the simulated period within each run and then as the median across repetitions. The subsection then identifies the contract conditions most strongly associated with participation and compares charging policies across locations to assess where charging control is most clearly related to system performance.

Higher V2G participation was associated with lower congestion loss (Figure 5.13), although the relationship was not strong enough for participation alone to explain system performance. Across all configurations, the pooled Spearman correlation was $\rho = -0.272$ ($p < 0.001$), while the correlation calculated within learning round and model version was slightly stronger at $\rho = -0.308$. Median loss decreased most clearly when participation increased from low to medium-high levels. Above approximately 400 participants, however, the improvement levelled off and median loss increased slightly.

Considerable variation remained at every participation level. Configurations with high participation could still produce high loss, while low-loss configurations occurred across a broader range of participation levels. Participation therefore created potential V2G flexibility, but its system value also depended on when vehicles were connected, their available charging and discharging flexibility, and the charging-control settings applied to the fleet. The association also differed between the two independent learning chains. It was weak and statistically inconclusive in V2G 1 ($\rho = -0.109$, $p = 0.058$), but substantially stronger in V2G 2 ($\rho = -0.558$, $p < 0.001$). V2G_1 contained fewer configurations at low participation levels and showed greater dispersion around the high-participation range, illustrating that the observed relationship was sensitive to the configurations explored in each learning chain.

Higher V2G participation was associated with lower congestion loss in both seasons and both learning chains. In both chains, this relationship was stronger in winter than in summer, suggesting that participation was more important under winter conditions, when congestion pressure was also greater. However, the strength of the relationship differed substantially between the learning chains. It was considerably stronger in V2G 2 ($\rho = -0.74$ in winter and $\rho = -0.64$ in summer) than in V2G 1 ($\rho = -0.42$ and $\rho = -0.35$, respectively). At comparable participation levels, absolute congestion loss was gener-




Figure 5.13: Relationship between mean V2G participation and congestion loss V2G 1 and V2G 2. Each point is an evaluated parameter configuration; congestion loss is shown on a logarithmic scale.

ally lower in summer than in winter, indicating that participation alone did not account for the seasonal difference in congestion performance.

Mean V2G participation could be predicted accurately from the experimental inputs, with an Extra Trees out-of-bag R^2 of approximately 0.94. Guaranteed range was by far the strongest predictor, accounting for approximately 0.59 of the normalised feature importance. Minimum plug-in time followed at approximately 0.12, battery degradation at 0.07, and financial compensation at 0.05. Participation was therefore determined primarily by the contract conditions rather than by charging infrastructure, charging policies, or renewable-system settings.

High participation was defined as the highest 20% of configurations, corresponding to a threshold of approximately 400 participants out of a fleet of 500. The selected PRIM box described this group using a maximum guaranteed range of 135 km, a minimum plug-in time of 1–5 hours, and battery-degradation values of 1 or 2. It captured 77.2% of high-participation configurations with a density of 74.2% and a lift of approximately 3.69. Narrower boxes reached densities above 95% after also restricting financial compensation to 60–150 and subsequently 90–150 per month, but their coverage declined to approximately 0.63 and 0.58. Compensation therefore further refined the high-participation region, but was less decisive than guaranteed range, minimum plug-in time, and battery degradation.

Charging-policy effects differed strongly by location (Figure 5.14). The clearest relationship was found at residential locations, where smart charging was consistently associated with substantially lower loss in both learning chains. In the pooled data, median loss was 0.55 under smart charging, compared with 3.80 for net-aware and 4.22 for normal charging. Smart charging differed significantly from both alternatives, whereas normal and net-aware charging did not differ significantly from each other. Both were also strongly represented among the worst-performing configurations. The residential policy represents charging control at residential locations more broadly, as the home-policy setting also applied to public chargers at residential destinations.

Differences among workplace charging policies were much smaller. In the pooled data, median loss ranged only from 0.58 for net-aware charging to 0.64 for smart charging and 0.69 for normal charging. None of the pairwise differences was statistically significant, and the ordering varied slightly between the two model versions. This suggests that workplace charging policy had no robust independent association with congestion loss within the explored parameter space.




Figure 5.14: Distribution of log-transformed congestion loss by charging policy (smart, net-aware, normal) at work, home, and public locations, for model versions V2G 1 and V2G 2.

Public and commercial charging policies likewise showed no stable ordering across model versions. In V2G 1, normal charging had a lower median loss than smart charging, 0.54 compared with 0.99, and this difference was statistically significant. However, the pattern was not reproduced in V2G 2, where smart charging had the lowest median loss and none of the pairwise differences was significant. In the pooled analysis, the median losses were closely grouped between 0.61 and 0.66, with no significant differences between policies.

Overall, higher V2G participation was associated with lower congestion loss especially in winter, but participation alone did not distinguish favourable configurations and its marginal benefit decreased at high participation levels. Guaranteed range was the dominant predictor of participation, while minimum plug-in time and battery degradation provided additional differentiation. Charging-policy effects were strongly location-dependent: smart residential charging was consistently associated with lower loss, whereas workplace and public charging policies showed much weaker and less stable relationships.

5.3.6. Alternative analyses and method selection

Several additional analyses were used to assess parameter effects and compare alternative interpretation methods. Descriptive comparisons first examined minima, medians, means, and the share of well-performing configurations for each parameter value. These plots repeatedly highlighted high renewable coverage, smart home charging, and medium-to-high guaranteed range. However, several contract and infrastructure parameters showed irregular or non-monotonic patterns. Because the Bayesian sampler evaluated parameter values unequally and in different combinations, these univariate summaries could not distinguish parameter-specific associations from interactions or changes across learning rounds. They were therefore used mainly for exploration rather than as the final basis for interpretation.

SHAP analyses were subsequently applied to several surrogate models and target definitions, including raw, log-transformed, and within-round-relative loss. The main global rankings broadly confirmed the feature-scoring results, with renewable coverage, home charging policy, and guaranteed range appearing most consistently among the leading parameters. SHAP additionally described how parameter contributions varied across individual configurations, but this did not substantially alter the global interpretation. Given the strong and stable predictive performance of the Extra Trees analysis and its more direct global ranking, feature scoring was retained as the principal method for parameter relevance, while SHAP was used as a supporting robustness check.

Permutation importance and analyses in which individual parameters were removed were also used to

test whether model performance declined when information from a parameter was no longer available. These analyses broadly confirmed the same leading parameters, but did not provide a substantially different interpretation. The scores were also less stable for the smaller round-specific datasets. They were therefore used as confirmation rather than as a separate results framework.

Partial-dependence plots, observed parameter-value comparisons, rank correlations, and other tests were used to examine the direction and shape of individual relationships. These analyses supported lower loss under high renewable coverage and smart home charging, while several contract and infrastructure parameters showed weaker or non-linear patterns. Their interpretation remained limited because partial dependence averages over the other inputs, including combinations that were sparsely represented, while univariate tests ignore interactions. PRIM was therefore preferred for identifying the direction and joint conditions associated with low-loss outcomes.

Finally, pairwise heatmaps, enrichment analyses, and dimensional-stacking plots were used to explore interactions between parameter values. These showed that favourable outcomes were generally associated with combinations rather than with isolated settings. Smart home charging and high renewable coverage appeared repeatedly, whereas the effects of contract and infrastructure values depended more strongly on the surrounding configuration. Many parameter pairs nevertheless contained relatively few observations, and the resulting patterns were difficult to summarise consistently across rounds and learning chains. PRIM provided a more compact multivariate description of these regions and was therefore selected for the final regional analysis.

Overall, the alternative analyses broadly confirmed the same dominant parameter patterns, increasing confidence that the main findings were not specific to a single interpretation method. Extra Trees feature scoring was retained for global parameter relevance, with SHAP and other importance methods providing supporting robustness checks, while PRIM was used to identify the combinations of parameter values associated with low-loss outcomes.

6

Discussion

This chapter interprets the main findings in relation to the research questions, model assumptions, and existing literature. It first discusses how V2G participation, contract design, charging control, renewable generation, and infrastructure influenced congestion-related outcomes. It then reflects on the value and limitations of the inverse-modelling approach, followed by the main model and experimental limitations that affect the interpretation and generalisability of the results.

6.1. Interpretation of results

The discussion follows the pathway through which V2G contributed to congestion relief. It first considers how contract preferences and behavioural feedback shaped participation, then how participation translated into usable flexibility during congestion peaks. It subsequently examines the wider energy-system, charging-control, and infrastructure conditions that determined whether this flexibility reduced congestion, before discussing the temporal pattern of the remaining congestion and the interpretation of the selected loss measure.

Throughout this discussion, congestion relief refers specifically to reductions in positive net-load exceedance above the selected 600 kW threshold. The objective did not penalise negative net load or other dimensions of network performance. Low-loss configurations should therefore be interpreted as favourable for the selected import-peak objective rather than as generally optimal electricity-system configurations. This qualification is important for the interpretation of the parameter rankings and favourable regions discussed below, which could differ under a broader or bidirectional congestion objective.

6.1.1. Drivers of V2G participation

Contract design was the main observed determinant of participation. The participation pattern broadly reflected the latent preference classes embedded in the model, but the simulation shows how this preference structure translated into realised participation after interaction with mobility behaviour and behavioural feedback. Guaranteed range was the clearest contract-related factor in the participation analysis, while minimum plug-in time and battery degradation played smaller secondary roles and financial compensation remained comparatively weak. The following subsections interpret how these attributes shaped participation, compare their realised importance with that implied by the inherited preference structure, and consider how modelled behavioural feedback and participation assumptions may have shaped these relationships.

Guaranteed range

Guaranteed range (GR) was the clearest contract-related driver of participation. Its importance was partly expected from the inherited preference structure: population-weighted latent-class preferences assigned approximately 46% of contract-attribute importance to GR. In the realised participation analysis, GR accounted for about 71% of the feature importance among contract attributes. Although these measures are not directly comparable, the difference shows that GR became even more dominant relative to the other contract attributes in the simulated participation outcomes.

In the model, range anxiety may have reinforced the direct effect of GR on participation. Higher GR reduced the likelihood that vehicles reached very low SOC levels, thereby limiting increases in range anxiety. To illustrate the magnitude of this mechanism with a GR of 75 km, an average trip left SOC only about 2 percentage points above the range-anxiety threshold, while one low-SOC interval required roughly 8.3 hours of driving at 50% SOC to offset. Low-SOC experiences could therefore have relatively persistent effects on later participation decisions. GR therefore supported participation both directly through contract utility and indirectly by limiting low-SOC experiences: the contribution of the latter pathway was not quantified separately.

The magnitude of the GR effect is nevertheless conditional on how mobility security was represented. In the model, agents relied on a fixed guaranteed range as a permanent mobility reserve and could not specify a desired SOC for a particular departure, as is possible in emerging V2G contracts (Burger, 2025). Range-anxiety feedback would then arise mainly when mobility needs deviate from what was planned, this could substantially reduce the importance of the GR compared with the present model. At the same time, permanent minimum-SOC settings are also used in practice (Octopus Energy, 2026), indicating that a protected reserve remains relevant. A more realistic representation would distinguish between planned and unexpected mobility needs.

Minimum plug-in time

Minimum plug-in time (MP) was the second-clearest contract-related predictor of participation, although its effect was substantially smaller than that of guaranteed range. MP was the dominant contract attribute only for the flexibility-focused latent class, which represented 18% of the population. Its relative importance in the participation analysis closely matched that implied by the population-weighted latent-class preferences, suggesting that its realised importance remained broadly similar.

The realised MP effect may also have been influenced by the implementation of plug-in concern (PIC). Participants were evaluated against their contractual MP, whereas non-participants were evaluated against a fixed eight-hour reference. Participants with short MP could therefore satisfy their benchmark relatively easily and reduce PIC, while non-participants who failed to reach eight hours could become more sensitive to plug-in requirements and consequently less likely to enter a V2G contract. This asymmetry may have favoured continued participation under short MP and discouraged entry among non-participants; its contribution was not quantified separately.

Battery degradation

Battery degradation (BD) had a relatively small effect on participation. This was somewhat weaker than expected from the inherited preference structure, in which degradation was negatively valued across all latent classes but was never the dominant attribute. In the realised participation analysis, BD explained comparatively little of the overall variation. Nevertheless, the highest participation levels occurred mainly under relatively low degradation assumptions. BD therefore appears to have acted less as a general driver of participation than as a limiting condition for reaching the highest participation levels.

This does not imply that degradation concerns are unimportant for V2G acceptance. Stated-choice studies generally find a negative effect, although its relative importance varies considerably across studies (Bakhuis et al., 2025; Huang et al., 2021; Meijssen, 2019). Battery ageing is most the explicitly addressed contractual factor in the Dutch roadmap for bidirectional charging, where clear information to users is emphasised (Ministerie van Infrastructuur en Waterstaat, 2026). However, battery ageing occurs regardless of V2G use, while the additional degradation associated with bidirectional operation depends on operating conditions and battery management (Lehtola, 2025). The relatively limited effect found here in combination with existing literature may therefore suggest that minimising stated degradation may not need to be the primary lever for increasing participation.

Financial compensation

Financial compensation (FC) had the weakest effect on participation among the contract attributes. Although it was the dominant attribute only for the smallest finance-focused class, representing 10% of the population, its realised importance was also substantially lower than suggested by the population-weighted latent-class preferences. This difference was partly structural: negative individual FC coefficients were constrained to zero, leaving approximately 44% of agents insensitive to changes in com-

pensation. The weak simulated effect should therefore not be interpreted as evidence that financial compensation is generally unimportant for V2G participation.

Even at very high participation levels, higher compensation was not consistently preferred. This suggests that compensation may be a basic condition for participation, while higher payments do not necessarily translate into proportionally higher uptake. All current V2G propositions reviewed likewise include some financial benefit, even when this only covers charging costs (Burger, 2025; Vattenfall, 2026).

Interpretation and scope of participation model

Contract re-evaluation occurred weekly, whereas real V2G contracts would likely involve longer commitments, switching costs, equipment investments, and administrative barriers. Agents also did not retain persistent individual travel patterns, so participation decisions could not reflect stable differences in mobility and parking requirements, particularly relevant to GR and MP. Furthermore, the experimental design allowed favourable levels of all contract attributes to be combined in a single configuration. High compensation, high guaranteed range, low plug-in requirements, and low perceived degradation could therefore occur simultaneously without explicit commercial trade-offs. This may have contributed to the high participation observed in some runs and makes direct comparison with realistic contract portfolios difficult.

Finally, participation was generated within a predefined latent-class structure and a limited set of psychological feedbacks. The model did not fully represent trust, contract complexity, privacy concerns, battery warranties, leasing arrangements, user override preferences, or competition between aggregators. Social diffusion and peer feedback were also not analysed despite evidence from EV-adoption research that social influence can affect adoption behaviour (Chakraborty et al., 2022). The resulting participation patterns should therefore be interpreted as behaviour under the selected preference structure, not as a complete forecast of real-world V2G adoption.

Overall, the realised participation pattern broadly reflected the inherited stated-choice preference structure, but not all contract attributes translated equally into simulated participation. Guaranteed range became relatively more dominant, minimum plug-in time remained broadly in line with its expected importance, and financial compensation played a substantially smaller role. Battery degradation explained little of the overall variation, although high degradation levels appeared to limit very high participation. Behavioural feedback may have modified these relationships over time, but RA, PIC, and ESI were not analysed separately and their individual contributions cannot be determined directly. Broader system parameters had little predictive importance for participation, indicating that participation remained primarily shaped by the contract and behavioural structure of the model.

6.1.2. From participation to usable flexibility

Higher V2G participation was generally associated with lower congestion loss, but participant count alone did not determine system performance. The following discussion first considers why the congestion benefit of additional participation diminished, then how vehicle availability and stored energy determined realised flexibility during congestion peaks, and finally how contract conditions affected the flexibility available from participating vehicles.

Participation and congestion relief

Participation was beneficial, but maximising participation was not equivalent to maximising congestion relief. Higher participation was moderately associated with lower loss, yet substantial variation remained at every participation level and low-loss configurations occurred across a broad participation range. The largest improvement occurred between low and medium-to-high participation, after which the median benefit levelled off. This is broadly consistent with previous work showing that substantial V2G system value can remain under partial participation, although higher participation can still improve system performance (Sagaria et al., 2024a).

Part of this apparent levelling-off may result from the loss formulation itself. Because only exceedances above 600 kW were penalised and larger exceedances received disproportionate weight, early peak reductions could produce larger improvements in loss than equivalent reductions once peaks had approached the threshold. However, the model also provides a control-based explanation for diminishing returns. The discharge capacity requested by the controller was based on the severity of net-load exceedance, independently of the number of participants. Realised discharge was then limited by the

usable V2G capacity available. A medium-sized participant pool may therefore already have been sufficient to meet much of the requested response, after which additional participants mainly spread the discharge across more vehicles rather than further reducing the peak.

Low-loss configurations may also have benefited from a division of roles within the wider EV fleet. Smart charging of non-participants could already avoid adding demand during critical periods, whereas participants followed the V2G controller and could additionally discharge. This interaction may partly explain why more participation did not translate monotonically into lower loss. The separate contribution of this interaction was not isolated, so it does not identify an optimal participation share.

The value of participation was also season-dependent. Its relationship with congestion loss was stronger in winter, when residual demand and congestion were substantially higher. Summer congestion was generally much lower, and several configurations eliminated positive-threshold loss altogether. Notably, the zero-loss summer configurations occurred only at relatively high participation levels, suggesting that high participation was most relevant for eliminating the final residual peaks rather than for achieving generally low summer congestion.

The key distinction is therefore between participation as potential flexibility and flexibility that can actually be deployed. This shifts attention from how many drivers participate to whether participating vehicles are connected and hold sufficient usable energy when congestion occurs.

Realised flexibility during congestion peaks

Participant count alone did not capture how much V2G flexibility was actually deployable during congestion peaks. Realised flexibility also depended on whether participants were connected, how much energy they held above their mobility reserve, and how charging and discharging were controlled. Average V2G state of charge was strongly negatively associated with loss, consistent with greater stored-energy availability supporting congestion relief. Because SOC was averaged over the simulation horizon, it does not directly measure energy available at individual peak events, and high SOC may also partly result from less intensive discharging. It nevertheless indicates that the energy state of the participating fleet mattered alongside its size.

The strong importance of residential charging control further suggests that usable flexibility depended on when energy was charged and retained, not only on how many vehicles participated. Charger-availability parameters had much smaller and less consistent effects, indicating that within the explored ranges, nominal access to charging infrastructure was less limiting than the timing and control of charging.

Contractual flexibility

Guaranteed range (GR) was not only the strongest contract-related predictor of participation, but also one of the contract parameters most consistently associated with congestion loss. Medium-to-high GR values occurred repeatedly in the low-loss region. Its association with congestion performance likely reflects a balance between maintaining a sufficiently large and stable participant pool and preserving usable discharge capacity. Higher GR could support participation, but simultaneously reserved more battery energy for mobility and therefore reduced the energy available for V2G discharge. The final results did not reveal a clear trade-off at which this flexibility cost outweighed the participation benefit. Both 105 and 135 km occurred in low-loss configurations with no significant difference, and loss did not increase consistently with GR. Within the tested range, the participation benefit and reduction in per-vehicle flexibility therefore appear to have largely offset one another, although this balance may shift beyond the tested range. The 100-kWh battery may have weakened the technical constraint by leaving substantial discharge capacity above the mobility reserve. Bakhuis (2026) observed the trade-off more clearly with a 50 kWh baseline battery, where contract changes attracting additional participants produced only limited extra balancing benefits as usable flexibility per vehicle declined. The present research shows a higher GR reduced the flexibility available per vehicle, but this did not result in a consistent increase in congestion loss within the tested settings.

High minimum plug-in time (MP) values were generally absent from the low-loss region, but their effect on congestion was mainly indirect. The model MP did not require vehicles to remain connected for the contracted duration and instead affected participation through plug-in concern. The limited MP effect may indicate that ordinary travel and parking patterns already provided sufficient connection time

in many cases, reducing the additional value of stricter requirements. This interpretation remains dependent on the behavioural plug-in patterns represented in the model. Connection requirements are also used in practice, although they are not universal; for example, Octopus Energy's Power Pack requires users to be able to remain plugged in for 12 hours per day on 20 days per month (Octopus Energy, 2026). In contrast to this research, Yun et al. (2025) illustrate a related trade-off: stricter minimum connection requirements reduced participation for several driver classes, but increased simulated peak-shaving potential by increasing vehicle availability.

Battery degradation (BD) and financial compensation (FC) had no direct operational effect on usable flexibility in the model. BD affected participation but did not reduce battery capacity, efficiency, or cycle life, while FC affected participation without changing vehicle availability or discharge capability. Their limited importance for congestion loss is therefore consistent with both attributes acting mainly through the size of the participant pool rather than through the flexibility provided by individual participants. Contract-related parameters were somewhat more important for raw than for log-transformed loss, suggesting that contract design played a larger role in distinguishing configurations with particularly high congestion loss than in explaining variation across the broader performance range.

Overall, participation created potential V2G flexibility, but congestion relief depended on how much of that flexibility was actually available when it was needed. Additional participation had diminishing value once the connected participant pool could provide much of the response requested by the controller, although participation was more strongly associated with lower loss in winter. State of charge, charging control, vehicle availability, and contract constraints therefore mattered alongside participant numbers. Guaranteed range was especially important because it supported participation while simultaneously reserving energy for mobility, whereas MP, BD, and FC affected usable flexibility more indirectly under the implemented model. Effective V2G therefore depended less on maximising participant numbers than on maintaining a sufficiently large pool of connected vehicles with usable energy during critical periods.

6.1.3. System conditions for congestion relief

The following discussion examines why low congestion loss was associated with a combination of higher renewable coverage, smart charging at residential locations, and a relatively solar-oriented generation mix, while charging infrastructure had a weaker and more conditional role.

Energy system

Renewable coverage was one of the clearest conditions associated with low congestion loss. Because net load was calculated as local demand minus solar and wind generation, greater renewable production directly reduced the residual demand that could exceed the congestion threshold. Charging control and V2G could reinforce this effect by shifting charging towards favourable periods and retaining energy for later discharge. The strong coverage effect is therefore partly inherent to the model structure, which contains no local fossil or other dispatchable generation source.

A relatively solar-oriented generation mix also occurred more frequently in the low-loss region. A plausible explanation is that daytime PV created charging opportunities before the evening peak because flexible charging increased only with sufficient headroom. Vehicles could absorb part of this production, enter the evening with a higher SOC, and require less charging when demand was already high. Some of the stored energy could subsequently be returned through V2G. The largest remaining losses nevertheless occurred during winter evenings, when solar production was absent or limited. The solar association should therefore be interpreted as an indirect interaction with charging control and total renewable coverage, rather than as evidence that solar generation directly supplied the most critical periods or generally outperformed wind.

The renewable result is conditional on how system performance was defined. Bakhuis (2026) found that higher renewable penetration could increase both balancing opportunities and the magnitude of imbalances, partly because wind-rich night charging could reduce battery headroom for later PV surpluses. The present model used a soft 80% SOC target to retain headroom, but production surpluses were not penalised by the congestion objective. The identified renewable configuration should therefore be interpreted as favourable for reducing positive net-load peaks, rather than for balancing imports and exports in both directions.

Charging control

Charging control across the wider EV fleet was one of the clearest conditions associated with congestion performance. Residential-location charging policy, which applied to both home and public chargers in residential areas, was one of the strongest predictors of congestion loss. Smart charging at these locations was repeatedly associated with low loss, whereas normal and net-aware charging were more strongly represented among the worst-performing configurations. This does not imply that smart charging could substitute for V2G, but it indicates that increasing V2G participation while leaving a substantial share of charging demand uncontrolled may leave an important source of congestion unresolved.

The stronger performance of smart charging is consistent with its ability to respond to realised net load and postpone demand, unlike normal charging and the predetermined net-aware profile. This likely complemented V2G by preventing additional EV demand during peaks, although their separate contributions were not isolated. Bakhuis (2026) similarly found that uncontrolled charging by non-participants remained an important source of imbalance.

The particularly strong residential effect is consistent with the timing of the modelled congestion peaks. Many vehicles returned to residential locations during the evening, when electricity demand was already high, and generally remained parked there for relatively long periods. This created both a risk and an opportunity: uncontrolled charging could reinforce the evening peak, while smart charging had sufficient time to postpone and spread demand. This residential focus also aligns with the Dutch roadmap for bidirectional charging, which identifies home charging as one of the first applications during the initial deployment phase (Ministerie van Infrastructuur en Waterstaat, 2026).

Charging-policy effects at workplace and public or commercial locations were substantially smaller and less consistent, with no robust pooled ordering across the two chains. This may partly reflect lower charging volumes and weaker alignment with the winter-evening peaks that dominated congestion loss. At work and commercial locations, net-aware charging differed little from normal charging during much of the day, while smart charging was constrained by the remaining parking duration. The predefined net-aware reductions may not always have been sufficient to suppress the full evening peak, while postponed charging could reappear later and shift rather than eliminate part of the peak creating the unstable results at residential locations. These explanations remain tentative because charging-policy combinations were sampled unevenly across the learning rounds.

Erden et al. (2019) similarly show that a mixed fleet with equal shares of smart-charging and V2G vehicles can fully shave the peak period when overall EV penetration is sufficiently high and uncontrolled charging remains limited. This illustrates that demand shifting by part of the fleet can complement discharge from another part once sufficient aggregate V2G capacity is available.

The charging and V2G control strategies were simplified rule-based mechanisms that responded to current net load and vehicle conditions without forecasting demand, renewable production, trips, or charger availability. The sequencing of charging, discharging, mobility, and net-load updates could therefore affect the realised response. The controller also responded gradually and did not necessarily use all technically available flexibility. Because charging and discharging were not jointly optimised across vehicles and future time steps, the identified configurations represent performance under the implemented control logic rather than the maximum technically achievable congestion relief. Forecast-based or jointly optimised control could therefore change the relative value of participation, SOC, infrastructure, and charging policies.

The charging-policy implementation was also stylised. Net-aware charging followed predetermined time-dependent limits, whereas smart charging responded to realised net load. Net-aware charging could therefore restrict charging when congestion was absent while missing peaks outside the predefined periods, while smart charging responded to realised net load but remained constrained by a fixed 80% SOC target. A lower or trip-specific SOC requirement could provide additional charging flexibility. The results therefore apply to these specific implementations rather than to all forms of grid-aware or smart charging.

Charging infrastructure

Charging-infrastructure parameters had weaker and non-monotonic effects, suggesting that their value depended on how the additional connections were used. Higher private or residential charger availabil-

ity appeared most useful when combined with smart charging, because more parked vehicles could access a charger at locations where they were likely to remain parked for longer. This gave the controller more opportunity to spread charging and maintain vehicles as potential V2G resources. Without coordinated charging, however, the same increase in access could also allow more vehicles to charge simultaneously and add to peak demand. Infrastructure therefore functioned mainly as an enabling condition whose congestion benefit depended on charging policy, location, and vehicle availability. Bakhuis (2026) similarly found that expanding charging infrastructure alone produced limited and ambiguous improvements in system balance. The limited sensitivity observed here may also partly reflect the simplified charging-power representation, with a uniform 10 kW limit across locations. More connected capacity therefore increased technical opportunity, but its realised congestion value still depended on how and when that capacity was used.

Overall, low congestion loss depended on the interaction between the energy system and how EV flexibility was controlled. Higher renewable coverage reduced positive net load, while residential smart charging helped prevent EV charging from reinforcing congestion peaks. A relatively solar-oriented generation mix was also associated with low loss, whereas charging infrastructure played a more conditional role. Effective congestion relief therefore depended less on maximising any single resource than on coordinating renewable availability, charging demand, and deployable EV flexibility in time.

6.1.4. Temporal patterns and residual congestion

The temporal analysis showed that congestion remained concentrated in a limited number of recurring winter-evening periods, even in better-performing configurations. The following discussion first considers when and why residual congestion persisted, and then examines what the selected congestion-loss measure captures and how its aggregated representation limits interpretation as physical grid congestion.

Temporal concentration

Congestion remained concentrated in winter evenings across both better- and worse-performing configurations. Lower-loss configurations reduced the frequency and magnitude of exceedances but did not fundamentally shift their seasonal or daily timing. This suggests that demand, renewable production, and mobility patterns primarily determined when congestion occurred, while charging and V2G settings mainly affected its severity. The interventions therefore reduced the existing peak problem rather than shifting it to another period.

The persistence of an evening residual peak is consistent with previous studies. Erden et al. (2019) found that a residential evening peak could only be shaved over the full peak period once sufficient connected V2G energy was available. Meyer et al. (2022) similarly found residual evening peaks whose reduction depended on control strategy and local system conditions. A smaller morning contribution also remained in better-performing runs, possibly reflecting rising service-sector demand and charging after the morning commute before solar production reached its daytime maximum. This may also indicate that the best-performing configurations redistributed charging demand away from the dominant evening peak without fully optimising its timing across the day.

Congestion loss was also highly concentrated in a relatively small number of five-minute intervals. Better configurations removed many small and medium-sized exceedances, leaving a limited number of persistent extreme events. This concentration reflects both the temporal structure of congestion and the squared loss function, which gives greater weight to large exceedances. The metric therefore distinguishes severe peaks from frequent marginal threshold crossings, but also makes the ranking particularly sensitive to a small number of extreme intervals. Many smaller and intermediate congestion periods could therefore be reduced through charging control and V2G, while the remaining extreme events may require more anticipatory control or complementary flexibility.

Interpretation of congestion loss

Low loss was not equivalent to a flat or fully balanced net-load profile. The objective penalised only positive net load above the selected threshold, while negative net-load valleys made no contribution. A configuration could therefore reduce winter demand peaks while creating deeper renewable surpluses at other times, and could even have a wider overall net-load range than a higher-loss configuration. Low loss should consequently be interpreted as successful reduction of the selected positive congestion

indicator, rather than as minimum total volatility or broader electricity-system performance. This also qualifies the relatively solar-oriented low-loss region and the strong association with high renewable coverage: these conditions were favourable for the selected positive net-load objective, while their effects on negative net load were not included in the ranking. The identified parameter importance and favourable regions were therefore partly conditional on this one-sided definition of congestion and could have differed under a broader objective.

The one-sided objective was a deliberate choice to focus on demand-side capacity constraints, where high electricity demand places pressure on the available grid capacity. This is particularly relevant for further electrification, as insufficient capacity during demand peaks can restrict additional electricity use (International Energy Agency, 2023b). The objective therefore captured the potential of V2G to create additional headroom during high-demand periods, rather than attempting to optimise both directions of net load simultaneously.

A more fundamental limitation concerns the representation of congestion itself. Congestion was approximated through aggregated positive net-load exceedances above a fixed threshold rather than through explicit network constraints. The model did not represent feeder loading, transformer capacity, voltage limits, power flows, or the spatial routing of electricity. The electricity system was effectively represented as a single aggregated balance rather than as a calibrated low- or medium-voltage distribution network. The fixed loss threshold was therefore not linked to the rated capacity of a specific feeder, transformer, or substation, and the model could not identify where within the network an exceedance would occur. It also could not represent situations in which one local branch is congested while sufficient capacity remains elsewhere. Low loss therefore indicates reduced aggregate positive peak pressure under the selected capacity proxy, but not the confirmed absence of distribution-grid congestion. The congestion measure should consequently be interpreted as a stylised indicator of local capacity pressure rather than as a representation of physical network feasibility.

6.2. Model scope, uncertainty and generalisability

Generalisability is limited by uncertainty in both the external conditions represented in the simulations and the chosen representation of the V2G system. The temporal inputs introduce external-factor uncertainty because mobility, residential and service-sector demand, and renewable production were derived from historical or stylised datasets that did not originate from the same year. Their combination may therefore create alignments between travel, demand, and renewable supply that did not occur historically. This is particularly relevant because congestion depended strongly on the coincidence of vehicle availability, charging demand, and renewable production. Mobility and electricity-use patterns may also have changed since the underlying datasets were collected, for example through increased working from home, electrification, efficiency improvements, and changing occupancy patterns. Moreover, a single weather and demand realisation cannot represent the range of possible years. The identified configurations should therefore be tested under alternative weather years, demand profiles, and mobility patterns to assess their robustness to external conditions.

Smart charging and V2G were controlled in response to modelled net load rather than electricity prices or market incentives. In practice, price signals reflect broader supply, demand, and market conditions and need not coincide with local grid constraints. The results therefore show the congestion-reduction potential of directly grid-responsive control, but not whether price-based operation would produce the same charging and discharging patterns.

The simulation horizon further limits temporal generalisability. The final runs did not cover a complete calendar year, leaving part of the final winter period unrepresented even though the largest congestion losses occurred in winter. This may have affected the relative contribution of seasonal events and consequently the low-loss regions identified by the sampler. Behavioural adaptation and seasonality also occurred simultaneously, as participation and psychological states continued to develop while external conditions changed. A multi-year simulation could use an initial stabilisation period followed by a complete seasonal evaluation under more settled behavioural conditions. The current horizon was sufficient to reveal broad seasonal and behavioural patterns, but not to establish their robustness across complete years or different weather conditions.

6.3. Interpretation, value and limitations of the approach

This section reflects on the methodological performance, interpretability, and limitations of the inverse-modelling approach. It first discusses sampler learning, then the consequences of the sampling design for inference and robustness, and finally positions the approach relative to related methods and its potential applications.

6.3.1. Search performance and surrogate learning

The inverse search identified favourable regions early, primarily by reducing the frequency of clearly poor proposals rather than by continuously improving the best configurations. This indicates that Bayesian Ridge was useful for directing the search away from clearly unfavourable configurations and towards lower-loss parts of the parameter space, which was the main purpose of the surrogate.

Bayesian Ridge represented parameter effects relatively simply, while the ABM could contain non-linear relationships and interactions between parameters. This simplification was particularly relevant for nominal charging-policy variables, which were encoded by their position in the parameter space before fitting Bayesian Ridge. The surrogate therefore treated these categories as ordered and equally spaced, although no such metric relationship was intended, which may have affected their predicted effects and candidate ranking. A more flexible surrogate, such as a tree-based model, could potentially have represented non-linearities and interactions more effectively.

Because the sampler was not compared with a random or space-filling design under the same simulation budget, its efficiency relative to alternative sampling strategies cannot be established directly. Surrogate accuracy was also not evaluated on an independent set of configurations during the search, so its value as a screening model is inferred mainly from the improvement in sampled outcomes rather than from independent predictive validation.

The surrogate was trained on individual repetitions rather than configuration-level summaries. Occasional high-loss repetitions could therefore influence learning more strongly than the median-based performance used in the subsequent analysis, and configurations with more observations received correspondingly greater weight. Batch sizes varied slightly across some of the experiment series, so summary statistics and contributions to subsequent surrogate updates were not perfectly balanced. However, these differences were small and are unlikely to have materially affected the overall learning pattern.

The later plateau suggests that additional rounds provided progressively less new information for distinguishing configurations within the already favourable region. As the search progressed, proposals concentrated within a narrower more favourable part of the parameter space, improving efficiency but reducing information about alternative regions. The formal stopping criterion was not reached, but it appears to have been conservative for the present purpose, as occasional high-loss proposals continued to affect the upper quartile after the favourable region had largely stabilised. The two search streams both moved towards lower and more concentrated loss distributions, which provides some support for the reproducibility of the broad search direction, but not demonstrating that all relevant favourable regions had been identified. Overall, the sampler was most successful at identifying and concentrating on promising regions, rather than at resolving small performance differences within those regions.

6.3.2. Consequences for inference and robustness

Adaptive sampling

Adaptive sampling produced an intentionally uneven sample, with progressively more evaluations in lower-loss parts of the parameter space. This increased the density of observations available for characterising favourable configurations, but reduced coverage of the broader parameter space. The resulting dataset was therefore well suited to identifying and interpreting candidate design regions, but less suited to global sensitivity analysis or balanced comparisons across the complete parameter space.

This sampling structure also affects how the subsequent parameter analyses should be interpreted. Associations, categorical comparisons, and feature importance scores describe variation within the configurations explored by the sampler rather than balanced parameter effects across the original parameter space. Because the same adaptive dataset was used both to locate favourable regions and to examine what characterised them, these interpretations remained exploratory rather than independently

confirmed. A separate balanced experiment would be needed to test whether the same relationships also hold more broadly outside the regions emphasised by the sampler.

Adaptive sampling can additionally change the apparent importance of parameters over time. Once the sampler repeatedly selects a favourable range of a parameter, that parameter varies less in the later data and may consequently receive a lower feature importance, even if it was important for reaching the favourable region in the first place. Other parameters may then appear more important because they explain the remaining variation within that narrower region. Feature importance should therefore be interpreted as context-dependent within the evolving sampled distribution.

The post-hoc analysis methods were nevertheless useful for characterising the sampled space. Extra Trees provided a relatively stable and interpretable ranking of predictive parameter relevance within the adaptively sampled configurations, rather than a global sensitivity measure over the complete parameter space. PRIM translated favourable outcomes into joint parameter restrictions. PRIM's axis-aligned boxes are useful for communicating compact regions but can underrepresent disconnected or more complex regions, which is particularly relevant if several qualitatively different configurations achieve similar performance. The analysis therefore cannot establish that all forms of equifinality were identified. At the same time, SHAP, permutation importance, partial dependence, and descriptive analyses broadly recovered the same dominant parameter patterns as the principal analyses. Although these methods were applied to the same adaptive dataset and therefore do not provide independent validation, their agreement supports the robustness of the main interpretation within the sampled ensemble.

Loss definition

The selected squared threshold loss made the search trajectory inherently objective-specific: only the magnitude of positive net-load exceedances above 600 kW contributed to sampler learning. The methodological system-level meaning of this objective was discussed earlier, the identified regions should be understood as favourable under the objective that guided the search rather than as generally optimal V2G configurations.

The shape of the loss also affected learning and interpretation. The threshold created a sharp distinction between penalised and unpenalised time steps, so small stochastic differences around 600 kW could change the loss, while the quadratic penalty amplified differences at the upper end of the distribution. This produced a strongly right-skewed response that was difficult for the linear surrogate to approximate precisely. The large separation among poor configurations relative to the narrower range among favourable configurations may help explain why the sampler was more effective at coarse screening than at resolving small differences within the low-loss region. The logarithmic transformation used in subsequent analyses compressed these extremes and improved predictive fit, but shifted the interpretation from absolute towards more proportional differences in loss.

Alternative loss specifications were evaluated post hoc on the same simulated runs to assess how sensitive configuration rankings and repetition-level stability were to the loss definition. None of these alternatives, however, guided its own adaptive search, so a different objective could have changed which configurations were sampled and consequently which favourable regions were identified. The strong overlap in rankings across several plausible loss formulations nevertheless suggests that the broad ordering of configurations was not highly sensitive to the exact loss specification.

Robustness

Ten repetitions were sufficient to obtain reasonably stable median rankings, but provided less reliable estimates of between-run variability and infrequent extreme outcomes. MAD should therefore be interpreted more cautiously than median performance. Importantly, lower median loss generally coincided with lower rather than higher MAD, providing no indication that lower-loss configurations were systematically less stable across repetitions. However, this pattern was loss-dependent and was not reproduced by the exceedance-count metric.

The stability series were not fully representative of the entire explored parameter space. Four of the six series focused mainly on regions that had already produced relatively low loss, while the two earlier series also included poorer-performing configurations. The estimated stability may therefore be somewhat more favourable than across a more balanced sample of the parameter space. However, the main relationship between repetition count and ranking stability was similar across both types of

series, supporting the use of ten repetitions for central comparisons while leaving stronger claims about dispersion and rare outcomes less certain.

Robustness was therefore assessed at several levels, but to different degrees. Stochastic robustness was examined most directly through repeated simulations, while the two independent search streams provided some evidence of consistency in the broad search direction and the alternative loss formulations showed that the broad configuration ranking was relatively insensitive to several plausible objective formulations. These checks support the robustness of the main performance patterns within the simulated setting, but do not establish robustness to alternative external conditions or model structures, which were discussed in Section 6.2.

6.3.3. Positioning and applicability of the approach

As positioned in Section 2.4, the present approach can be understood as a performance-guided, future-oriented form of inverse exploration within a largely predefined ABM. The results now allow this positioning to be assessed more specifically in relation to optimisation and exploratory modelling. The purpose here is not to assign the approach to a single methodological category, but to compare the methods based on characteristics.

Relation to related methods

Operationally, the procedure shared important characteristics with optimisation. The present study may be more closely related to optimisation than explanatory inverse-modelling applications are, because congestion-related loss defined an explicitly desirable performance direction and many of the searched inputs represented possible system or design choices. This resemblance was strengthened by the decision to use purely exploitative sampling in the final learning phase, which directed new evaluations towards configurations predicted to have lower loss rather than balancing predicted performance with uncertainty. The analytical endpoint differed from conventional optimisation, however, because the aim was not to identify or refine a single optimum but to interpret broader combinations of conditions associated with favourable performance. A potential advantage of examining a broader favourable region rather than relying on one numerically best configuration is that the interpretation becomes less dependent on the precise identity of a single best-performing solution. This is consistent with Epstein (2023), who emphasises that inverse search may yield multiple comparably fit generators and that common elements across such solutions can themselves be informative.

The approach also overlaps with Exploratory Modelling and Analysis (EMA). In its more conventional forward application, EMA distinguishes controllable policy levers from uncertain external conditions and evaluates how alternative policies perform across plausible futures (Bankes, 1993; Kwakkel et al., 2016). The inverse formulation used here instead started from a desired performance direction and searched for combinations of inputs associated with that outcome. A useful difference in emphasis may be evaluating the robustness of policies across uncertain futures and searching backwards from a target to identify conditions or pathways capable of producing favourable performance.

This distinction is not absolute. EMA can also use directed search to identify combinations associated with outcomes of interest, while an inverse search could evaluate candidate configurations across multiple scenarios or uncertainties. One potentially useful difference in implementation lies in how the search is guided. A Bayesian inverse search can use surrogate predictive uncertainty alongside predicted performance, allowing evaluations to be directed both towards the target and towards poorly understood parts of the parameter space. This can support characterisation of a broader favourable region rather than only refinement of the currently best-predicted solution. This uncertainty concerns uncertainty in the surrogate representation of the response surface and is distinct from the uncertain future conditions commonly explored in EMA. The distinction is nevertheless not inherent to either framework, since similar acquisition strategies could also be incorporated into EMA. In the present study, this potential was only partly used because the final learning phase prioritised exploitation of predicted low loss.

Applicability and limitations

A further implication concerns the definition of success. Unlike explanatory inverse modelling, where the target is generally derived from an observed pattern, a future-oriented target is selected by the modeller and is therefore partly normative. The chosen loss function consequently determines not only

how outcomes are evaluated but also where simulation effort is concentrated. This makes forward-model credibility and the completeness of the performance criteria particularly important. Where design problems contain important trade-offs, several objectives or complementary outcome constraints may be preferable to an increasingly precise search around one narrowly defined performance measure. Rather than removing equifinality, such an approach could help reveal alternative acceptable solution directions and their trade-offs.

Overall, performance-guided inverse exploration appears particularly useful when the configuration space is large, the desired outcome can be defined sufficiently clearly, and the research objective is to discover and understand several satisfactory configurations or favourable regions rather than one precise optimum. Inverse modelling can also be appropriate when uncertainty makes the identity of a single optimum unstable or not especially meaningful. Potential applications include early-stage system and policy design, where multiple pathways should remain open before path-dependent investments are made; screening large configuration spaces before more detailed robustness analysis; and behavioural or other equifinal models in which several mechanisms or parameter combinations can generate similar outcomes. In such applications, the value lies not only in locating favourable solutions, but also in identifying recurring conditions and alternative pathways across them.

6.4. Recommendations for future research

First, future research should embed the charging and V2G model in a more explicit distribution-network representation. Spatially resolved bidirectional power flows and relevant network constraints would allow congestion effects to be evaluated more directly than through aggregate net load alone. This would also make it possible to assess whether periods of negative net load create local export-related constraints, which were not captured by the present loss function. The identified low-loss configurations could then be tested for whether they relieve congestion at the network locations where it actually occurs.

Second, the current rule-based charging and V2G controller should be compared with anticipatory alternatives. Forecast-based and multi-period optimisation could prepare vehicles for expected winter peaks, preserve sufficient state of charge before long events, and coordinate charging and discharge across time. Price-oriented control should also be tested separately, because a strategy that performs well in energy markets does not necessarily reduce local congestion. The order in which charging and discharge rules are applied should form part of this comparison. The congestion benefits of smart charging and V2G should also be compared with their implementation and operational costs. This would allow future work to assess which combinations of charging control, V2G participation, infrastructure, and contract conditions provide congestion relief most cost-effectively, and whether the additional flexibility provided by V2G justifies its additional costs relative to smart charging alone.

Third, simulations should cover multiple consecutive years. At least one initial year could be used for behavioural stabilisation, followed by two or more evaluation years with different weather, demand, and mobility profiles. Input data should preferably be temporally aligned and mobility assumptions should reflect more recent post-pandemic travel behaviour. Persistent driver-specific mobility schedules would provide a more realistic basis for psychological feedback, while allowing routine and incidental trips to be distinguished and range anxiety to reflect expected mobility needs rather than a fixed daily reserve. Private home charging and public charging at residential destinations should be represented and analysed separately.

Fourth, the inverse-modelling procedure and the subsequent analysis methods should be evaluated more systematically. Bayesian-guided sampling (explorative and exploitative) should be compared with random, space-filling, and alternative adaptive strategies under an equal computational budget to determine when directed sampling provides a meaningful efficiency advantage. A separate balanced validation sample could then test whether the dominant parameter relationships and favourable regions identified from the adaptive search persist under more even coverage of the parameter space. Together, these tests would show whether the efficiency gained through directed sampling outweighs the loss of global coverage and whether conclusions drawn from the resulting adaptive dataset remain stable under more balanced sampling.

Finally, future research should also compare the data requirements and stability of alternative interpre-

tation methods under different sampling conditions. PRIM, feature scoring, SHAP, partial dependence, and related scenario-discovery techniques could be applied to progressively smaller samples drawn from both balanced and adaptively concentrated datasets. This could establish which methods provide stable and informative results for different types of simulation data, and how much data each requires, potentially reducing the number of computationally expensive ABM runs needed for interpretation.

7

Conclusion

This research investigated how Vehicle-to-Grid systems can be designed to mitigate local grid congestion when user participation is not imposed as a fixed assumption but emerges from heterogeneous contract preferences, mobility behaviour, and psychological feedback. The study combined a literature review with a full-year agent-based simulation model and a learning-guided inverse-modelling procedure. The model connected contract acceptance, mobility, charging opportunities, charging control, renewable generation, and V2G operation to aggregate local net-load outcomes. Congestion performance was evaluated through positive net-load exceedances above a selected threshold, while repeated simulations and two independent learning streams were used to account for stochastic variation and dependence on the sampled search trajectory.

The aim was not to identify one technically optimal V2G system. Instead, the study sought to determine which combinations of behavioural, contractual, infrastructural, and operational conditions repeatedly produced lower congestion-related loss, and whether inverse modelling could be used identify these design regions without exhaustively evaluating the complete parameter space.

7.1. Main research question

The main research question was:

Which Vehicle-to-Grid (V2G) system configurations mitigate local grid congestion in Dutch urban areas when user participation is modelled endogenously?

Within the stylised urban system and for the selected objective of reducing positive net-load exceedances, favourable congestion performance was most consistently associated with a combination of smart charging at residential locations, high renewable-energy coverage, and a medium-to-high guaranteed range. A relatively solar-oriented generation mix, avoidance of the most demanding minimum plug-in requirement and highest battery-degradation value, and sufficient private-charging availability further refined this low-loss region. These conditions did not operate independently. They formed a system in which residual electricity demand, user participation, vehicle availability, and charging control were better aligned around the periods in which peak pressure occurred.

Participation was an important enabler of congestion reduction because it created a pool of batteries that could potentially be discharged during high-load periods. However, increasing participation did not produce proportional improvements across the complete range. The largest reduction in loss occurred between low and medium-high participation, after which the improvement levelled off. Consequently, the best-performing systems did not simply maximise acceptance. They ensured that a sufficiently large group of participants was connected, had usable energy available above its mobility reserve, and could be controlled when congestion occurred. Guaranteed range was important in this respect because it determined how much battery capacity remained reserved for mobility and therefore how much flexibility was available for V2G discharge.

The most effective configurations combined two complementary mechanisms. Smart charging reduced or postponed EV demand that would otherwise have entered an existing peak. V2G subsequently addressed part of the positive residual net load that could not be avoided through charging shifts alone. Effective V2G design therefore requires coordination of the charging demand of the wider EV fleet, including non-participants, rather than only increasing the number of bidirectional vehicles.

7.2. Contractual and psychological determinants of participation

The first sub-question asked:

How do psychological and contractual factors influence EV owners' willingness to participate in V2G systems?

The literature shows that V2G participation depends not only on financial remuneration, but also on mobility security, autonomy, and flexibility. Guaranteed range, minimum plug-in requirements, expected battery degradation, and compensation all influence acceptance, although their relative importance differs between user groups. Psychological concerns such as range anxiety and perceived loss of control cannot necessarily be offset by higher payments. Participation is therefore heterogeneous, and one contract is unlikely to be equally attractive to all EV users.

The simulation results reinforced the importance of mobility security. Guaranteed mobility security played a central role in shaping participation outcomes, followed at a considerable distance by minimum plug-in time. Battery degradation was less influential overall, although high degradation still limited the highest participation levels, while financial compensation had only a weak influence, partly because not all modelled drivers were sensitive to changes in compensation.

The influence of guaranteed range may have been reinforced by behavioural feedback. A higher reserve increased contract utility while also reducing low-state-of-charge experiences that could increase range anxiety. This effect may be smaller when drivers retain greater control over departure timing or minimum SOC. Minimum plug-in time should be interpreted more cautiously. In the model, it affected utility and plug-in concern but did not guarantee that vehicles were physically connected for the specified period. Drivers did not follow fixed personal travel routines, which limited how strongly participation could reflect persistent differences in mobility.

Overall, participation reflected the underlying heterogeneity in user preferences. Contracts therefore attracted participation most effectively when they protected mobility needs, limited plug-in burdens, and reduced concerns such as battery degradation, rather than relying primarily on higher payments.

7.3. Participation and congestion relief

The second sub-question asked:

How does user participation affect the effectiveness of V2G systems in mitigating local grid congestion?

Higher participation was, on average, associated with lower congestion-related loss. The largest improvement occurred when participation increased from low to medium-high levels, while the additional benefit diminished at very high participation levels. Poor outcomes nevertheless still occurred at high participation levels, showing that participant count alone did not determine congestion relief.

This levelling-off was partly shaped by how discharge was controlled, additional vehicles added little congestion relief once enough vehicles were available to provide the required response to a peak. The way congestion was evaluated also contributed to this pattern, because the highest peaks were weighted more heavily. High participation was the most important in winter, when congestion was highest. In summer, the remaining modelled congestion pressure was eliminated only at very high participation levels.

Participation therefore represented potential flexibility rather than direct system value. Its contribution depended on whether vehicles were connected during critical periods, had sufficient energy available above their mobility reserve, and could be coordinated effectively. In the model, increasing V2G par-

ticipation could also reduce the share of vehicles governed by smart-charging rules. Additional participation was therefore not necessarily beneficial when it replaced charging control that would otherwise have prevented demand from entering the peak.

Overall, sufficient participation remained essential for reducing loss up to a certain level, but maximising acceptance was not equivalent to optimising congestion reduction. A smaller, well-controlled and operationally available participant pool could provide more value than a larger pool with limited availability or weaker charging coordination.

7.4. Smart charging, grid-aware charging, and V2G

The third sub-question asked:

How do smart charging and grid-aware charging contribute to V2G in mitigating grid congestion under different system design configurations?

Smart charging at residential locations was one of the conditions most consistently associated with lower congestion pressure. This partly reflects that the residential policy applied to both private home chargers and public chargers at residential destinations, covering a large share of charging opportunities. Smart charging reduced avoidable EV demand during high-load periods, while V2G provided additional flexibility by discharging when positive net load remained high. The two strategies were therefore complementary. However, their separate contributions were not isolated, so the results do not show that either strategy was generally more effective than the other.

The findings for workplace and public or commercial charging were less clear, partly because charging volumes were lower, parking time was shorter and vehicle presence overlapped less with the severe winter-evening peaks. Grid-aware charging likewise showed no clear advantage over normal charging under the tested conditions, so its contribution to congestion mitigation could not be established clearly from the present model. Finally, all charging and V2G strategies were configured specifically to reduce congestion-related net load. Their relative performance should therefore not be transferred directly to controllers aimed at wholesale-price optimisation, energy-market revenues, or user electricity costs.

7.5. System-design conditions and the origin of congestion

The fourth sub-question asked:

Which combinations of contractual, infrastructural, and operational parameters define favourable congestion outcomes?

Across alternative definitions of favourable performance, the most stable design core consisted of smart charging at residential locations, high renewable-energy coverage, and a medium-to-high guaranteed range. As the low-loss definition became more selective, this core was refined by a relatively solar-oriented generation mix, avoidance of the highest minimum plug-in requirement and battery-degradation value, and sufficient private-charging availability. Private-charging availability appeared mainly relevant in combination with smart residential charging, while infrastructure effects were otherwise relatively weak. The analysis therefore identified nested favourable design regions rather than one exact configuration required for low congestion pressure.

These findings are conditional on the one-sided congestion objective, which penalised only positive net-load exceedances above 600 kW. Higher renewable coverage reduced the residual demand to which charging control and V2G had to respond, while the solar–wind mix influenced when energy was available relative to demand and vehicle availability.

Congestion remained concentrated in winter evenings in both better- and worse-performing configurations. Demand, mobility, and renewable-production profiles primarily determined when congestion occurred, while charging and V2G control mainly determined its severity. Better configurations removed many small and medium-sized exceedances without fundamentally shifting congestion to another season or time of day.

The remaining loss was therefore increasingly concentrated in a limited number of persistent extreme

peaks. These events required flexibility that was sufficient for both the magnitude and duration of the full peak and may ultimately require forecast-based control, multi-period optimisation, complementary flexibility or storage, and physical network reinforcement.

7.6. Methodological contribution of inverse modelling

The final sub-question asked:

How does Bayesian-guided inverse modelling support the search for promising V2G configurations within the large parameter space of a stochastic agent-based model?

The approach may be best described as a performance-guided inverse exploration of a largely predefined, behaviourally explicit ABM. The Bayesian Ridge-guided search was most useful for identifying and concentrating on favourable regions rather than for finding a precise optimum. The largest improvement occurred early, while later rounds mainly avoided poor configurations. This pattern may partly reflect the congestion-loss formulation, which strongly distinguished severe peaks, together with the exploitative search. Learning nevertheless remained non-monotonic, consistent with the difficulty of approximating the nonlinear and stochastic ABM with a relatively simple surrogate.

The adaptive search intentionally produced more information about promising regions than about the complete parameter space. This was useful for the purpose of the inverse search, but means that the resulting parameter relationships describe the adaptively explored configurations rather than the full system space. Region-based analysis subsequently translated the sampled results into interpretable combinations of favourable conditions rather than relying on individual best-performing configurations.

A central implication is that the inverse search remained conditional on the performance objective chosen by the modeller. The identified regions are therefore favourable for the selected congestion objective within the modelled system, rather than empirical predictions or generally optimal V2G designs. One of the main methodological values of the approach lies in narrowing a large configuration space towards interpretable promising regions that can subsequently be examined more thoroughly.

7.7. Implications for system design and policy

7.7.1. Aggregators and contract design

The implementation of flexible charging requires coordination between actors that contract EV users, identify local network needs, operate charging infrastructure, and control vehicle charging. These roles may be fulfilled by energy suppliers, aggregators, mobility service providers, charge-point operators, fleet operators, grid operators, or combinations of these parties. Effective coordination therefore requires clear agreements on responsibilities, control rights, and information exchange, supported by interoperable technical and communication standards.

Contract design should mainly focus on mobility security while avoiding unnecessary participation burdens. Long minimum plug-in requirements were not associated with lower congestion and may therefore be unnecessary when existing mobility patterns already provide sufficient connected flexibility. Contracts could instead focus on availability during specific congestion periods. However, this finding depends on the modelled mobility and connection patterns and should therefore not be interpreted as a general argument against minimum availability requirements.

High fixed compensation was not required for V2G itself and played only a limited role in encouraging participation. Where compensation is offered, it may therefore be more effective to link payments to the flexibility that is actually made available. This could encourage users to remain plugged in when flexibility is valuable without imposing restrictive minimum plug-in obligations. Users could also be rewarded for the amount of battery capacity they voluntarily make available while retaining control over their own minimum guaranteed range, thereby preserving mobility security and limiting range anxiety.

Contracting strategies should consequently focus on obtaining a reliable and operationally useful pool of flexibility rather than simply maximising participant numbers or offering high uniform payments. Combining V2G with smart charging can further reduce the residual peak before bidirectional discharge is required, lowering both the necessary V2G capacity and battery cycling. The aggregation strategy and

control signal remain important, because price-based steering alone may not alleviate congestion when market prices and local network peaks do not coincide.

7.7.2. Grid operation and charging-system

The location and timing of flexibility should be based on local demand, mobility, connection duration, and congestion patterns. Grid operators can identify where and when flexibility is required, while charging-system providers and other operational actors can translate these requirements into charging and discharging decisions. Residential destinations were the most relevant intervention point in the model, including both private home chargers and public chargers in residential areas. The results do not show that workplace charging has no value, but its contribution was less clearly identifiable under the modelled profiles. Infrastructure planning should therefore follow local demand, mobility, dwell-time, and congestion patterns.

Policy and implementation strategies should treat smart charging and V2G as complementary layers of flexibility. Smart charging can already reduce congestion pressure by shifting charging away from emerging peaks, while V2G adds a stronger response by discharging during the remaining critical periods. The low-loss configurations generally combined controlled charging with very high V2G participation, suggesting that the strongest congestion relief is achieved when both mechanisms are coordinated. Because their separate effects could not be isolated in the sampled configurations, the results do not show that smart charging alone is sufficient, but they do indicate that it provides an important foundation on which V2G can build.

To prevent flexible charging from reinforcing local peaks, price incentives and control signals should be aligned with local congestion conditions. Signals based only on wholesale electricity prices or renewable availability may otherwise encourage charging at times that are economically or environmentally favourable but unfavourable for the local network.

7.7.3. Public policy and system planning

Governments, regulators, and public authorities can support this system by shaping infrastructure deployment, market arrangements, interoperability requirements, data access, and network investment. V2G should be treated as a complement to grid reinforcement rather than as a substitute. Smart and bidirectional charging should receive particular attention at residential locations, where the model identified the clearest contribution, while workplace deployment should remain conditional on local system needs.

Policy should focus not only on the number of installed bidirectional chargers or signed contracts, but on whether they can provide useful flexibility during critical periods. Responsibilities for contracting users, identifying local flexibility needs, controlling vehicles, and delivering flexibility should therefore be clearly organised. Effective congestion relief also requires grid operators, aggregators, charge-point operators, and charging-system providers to exchange timely information about local net load, available flexibility, and realised response. Policy should facilitate this cooperation through interoperable communication standards, clear data-access arrangements, and appropriate safeguards for privacy and cybersecurity.

System planning should also retain other flexibility options, including stationary storage, building and industrial demand response, renewable curtailment, and network reinforcement, for the limited but severe winter peaks that EV flexibility could not resolve reliably.

7.8. Recommendations for future research

Future research should embed this framework into spatially explicit distribution networks to evaluate physical power flows and potential export congestion. Congestion assessment should also extend beyond positive net-load peaks by considering negative net load and, where relevant, multiple system objectives. The current rule-based controller should be compared with anticipatory, forecast-based, and price-oriented strategies while evaluating the cost-effectiveness of V2G relative to smart charging. Furthermore, multi-year simulations with updated and temporally aligned demand and mobility profiles are needed, preferably with more persistent driver-specific travel patterns to strengthen the representation of behavioural feedback and with private and public residential charging represented separately. Methodologically, Bayesian-guided inverse search should be benchmarked against alternative sam-

pling designs under equal computational budgets, supported by balanced validation samples, while the data requirements and stability of subsequent analysis methods should also be assessed.

7.9. Final conclusion

This study demonstrates that the contribution of V2G to congestion mitigation cannot be inferred from technical battery capacity or participation rates alone. Contracts were most effective at attracting participation when they protected mobility security and limited participation burdens. Sufficient participation was necessary to create an effective flexibility pool, but maximising acceptance was not equivalent to maximising congestion reduction. Actual system value depended on whether vehicles were connected, sufficiently charged, and controllable during the relevant peak.

Smart charging at residential locations was strongly complementary to V2G because it prevented avoidable charging demand from entering the dominant winter-evening peak. V2G could subsequently reduce part of the remaining positive net load. Within the explored system, the most favourable configurations combined this controlled residential charging with high renewable coverage and a medium-to-high guaranteed range, supplemented by a relatively solar-oriented mix, moderate contract burdens, and sufficient charging access.

Most congestion periods could be reduced through this combination, but a limited number of persistent extreme winter peaks remained. V2G should therefore be developed as one component of a coordinated flexibility system that combines user-compatible contracts, controlled charging, grid-responsive operation, complementary flexibility, and continued network reinforcement.

References

- Andersen, D., & Powell, S. (2025). Policy and pricing tools to incentivize distributed electric vehicle-to-grid charging control. *Energy Policy*, 198, 114496. <https://doi.org/10.1016/j.enpol.2025.114496>
- Autoriteit Consument & Markt. (2025, July). *Overzicht congestiemaatregelen acm* (tech. rep.) (Accessed 2026-06-18). Autoriteit Consument & Markt. <https://www.acm.nl/system/files/documents/overzicht-congestiemaatregelen-acm-juli-2025.pdf>
- Bakhuis, J. (2026). *Vehicle-to-grid in the netherlands: A sociotechnical multi-system analysis* [Dissertation]. Delft University of Technology. <https://doi.org/10.4233/uuid:9cd3b1a0-f04d-481b-a53e-2b60397bb9dc>
- Bakhuis, J., Barbour, N., Molin, E., & Chappin, É. J. (2025). Understanding user preferences regarding vehicle-to-grid (V2G): A latent class choice analysis. *Transportation Research Part A Policy and Practice*, 199, 104610. <https://doi.org/10.1016/j.tra.2025.104610>
- Banks, S. C. (1993). Exploratory modeling for policy analysis. *Operations Research*, 41(3), 435–449. <https://doi.org/10.1287/opre.41.3.435>
- Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2).
- Bibak, B., & Tekiner-Mogulkoc, H. (2021). Influences of vehicle to grid (V2G) on power grid: An analysis by considering associated stochastic parameters explicitly. *Sustainable Energy Grids and Networks*, 26, 100429. <https://doi.org/10.1016/j.segan.2020.100429>
- Broneske, G., & Wozabal, D. (2017). How Do Contract Parameters Influence the Economics of Vehicle-to-Grid? *Manufacturing And Service Operations Management*, 19(1), 150–164. <https://doi.org/10.1287/msom.2016.0601>
- Burger, J. (2025, September 11). *Comparing european v2g offers*. Retrieved August 20, 2026, from <https://www.watture.eu/comparing-european-v2g-offers/>
- CE Delft. (2023, October). *Beleid voor grootschalige batterijen en opweknetcongestie* (tech. rep.) (Accessed 2026-06-22). Rijksdienst voor Ondernemend Nederland. <https://www.rvo.nl/sites/default/files/2023-11/Rol-grootschalige-batterijen-bij-opweknetcongestie-CE-Delft.pdf>
- CE Delft. (2025, May). *Verkenning tijdsafhankelijke nettarieven voor grootverbruikers* (tech. rep.) (Accessed 2026-06-22). Netbeheer Nederland. https://www.netbeheernederland.nl/sites/default/files/2025-05/ce_delft_240472_achtergrondrapport_tijdsafhankelijke_nettarieven_grootverbruikers_def_002.pdf
- CE Delft & Merosch. (2024, June). *Oplossingen voor netcongestie bij bedrijven* (tech. rep.) (Accessed 2026-06-22). Rijksdienst voor Ondernemend Nederland. <https://www.rvo.nl/sites/default/files/2024-06/Rapport-RVO-CEdelft-Merosch-oplossingen-netcongestie-bedrijven.pdf>
- Centraal Bureau voor de Statistiek. (2016). *Onderzoek verplaatsingen in nederland (ovin)*. <https://www.cbs.nl/nl-nl/onze-diensten/methoden/onderzoeksomschrijvingen/korte-onderzoeksbeschrijvingen/onderzoek-verplaatsingen-in-nederland--ovin-->
- Chakraborty, D., Bunch, D. S., Brownstone, D., Xu, B., & Tal, G. (2022). Plug-in electric vehicle diffusion in california: Role of exposure to new technology at home and work. *Transportation Research Part A: Policy and Practice*, 156, 133–151. <https://doi.org/10.1016/j.tra.2021.12.005>
- Chen, C.-F., De Rubens, G. Z., Noel, L., Kester, J., & Sovacool, B. K. (2020). Assessing the socio-demographic, technical, economic and behavioral factors of Nordic electric vehicle adoption and the influence of vehicle-to-grid preferences. *Renewable and Sustainable Energy Reviews*, 121, 109692. <https://doi.org/10.1016/j.rser.2019.109692>
- Claessen, F. N., Claessens, B., Hommelberg, M. P. F., Molderink, A., Bakker, V., Toersche, H. A., & van den Broek, M. A. (2014). Comparative analysis of tertiary control systems for smart grids using the flex street model. *Renewable Energy*, 69, 260–270. <https://doi.org/10.1016/j.renene.2014.03.037>
- Dam, K. H., Nikolic, I., & Lukszo, Z. (2013). *Agent-Based Modelling of Socio-Technical Systems*. Springer. <https://doi.org/10.1007/978-94-007-4933-7>


- Davis, F. D., et al. (1989). Technology acceptance model: Tam. *AI-Suqri, MN, Al-Aufi, AS: Information seeking behavior and technology adoption*, 205(219), 5.
- De Graaf, J. A., Stok, F. M., De Wit, J. B., & Bal, M. (2026). Understanding behavioral factors influencing EV smart charging: A mixed-method study of citizens' capabilities, opportunities, and motivations. *Energy Policy*, 210, 115060. <https://doi.org/10.1016/j.enpol.2025.115060>
- Deru, M., Field, K., Studer, D., Benne, K., Griffith, B., Torcellini, P., Liu, B., Halverson, M., Winiarski, D., Rosenberg, M., Yazdanian, M., Huang, J., & Crawley, D. (2011). *U.S. Department of Energy Commercial Reference Building Models of the National Building Stock* (tech. rep.). National Renewable Energy Laboratory. <https://doi.org/10.2172/1009264>
- (DHPC), D. H. P. C. C. (2024). DelftBlue Supercomputer (Phase 2).
- ElaadNL & Vereniging Elektrische Rijders. (2025, September). *Nationaal laadonderzoek 2025* (tech. rep.) (Accessed 2026-06-22). ElaadNL. <https://elaad.nl/wp-content/uploads/downloads/Nationaal-Laadonderzoek-2025.pdf>
- ElaadNL, Team Marktontwikkeling. (2025, August). *Bidirectioneel laden: Welke kant gaat het op?* (Tech. rep.). ElaadNL. Arnhem. <https://elaad.nl/wp-content/uploads/2025/09/ElaadNL-Outlook-Bidirectioneel-laden.pdf>
- Elia. (2018). Wind-power generation data. <http://www.elia.be/en/grid-data/power-generation/wind-power>
- Eltohamy, M. S., Tawfiq, M. H., Ahmed, M., Alaas, Z., Mohammed, B., Ahmed, I., Youssef, H., & Raouf, A. (2025). A comprehensive review of Vehicle-to-Grid V2G technology: Technical, economic, regulatory, and social perspectives. *Energy Conversion and Management X*, 27, 101138. <https://doi.org/10.1016/j.ecmx.2025.101138>
- Epstein, J. M. (2023). Inverse generative social science: Backward to the future. *Journal of Artificial Societies and Social Simulation*, 26(2), 9. <https://doi.org/10.18564/jasss.5083>
- Erden, F., Kisacikoglu, M. C., & Erdogan, N. (2019). Adaptive v2g peak shaving and smart charging control for grid integration of pevs. *Electric Power Components and Systems*, 46(13), 1494–1508. <https://doi.org/10.1080/15325008.2018.1489435>
- Fagiolo, G., Windrum, P., & Moneta, A. (2007). A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. *Computational Economics*, 30(3), 195–226. <https://doi.org/10.1007/s10614-007-9104-4>
- Geske, J., & Schumann, D. (2018). Willing to participate in vehicle-to-grid (V2G)? Why not! *Energy Policy*, 120, 392–401. <https://doi.org/10.1016/j.enpol.2018.05.004>
- Ghotge, R., Nijssen, K. P., Annema, J. A., & Lukso, Z. (2022). Use before You Choose: What Do EV Drivers Think about V2G after Experiencing It? *Energies*, 15(13), 4907. <https://doi.org/10.3390/en15134907>
- Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., Thulke, H.-H., Weiner, J., Wiegand, T., & DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, 310(5750), 987–991. <https://doi.org/10.1126/science.1116681>
- Gschwendtner, C., Sinsel, S. R., & Stephan, A. (2021). Vehicle-to-X (V2X) implementation: An overview of predominate trial configurations and technical, social and regulatory challenges. *Renewable and Sustainable Energy Reviews*, 145, 110977. <https://doi.org/10.1016/j.rser.2021.110977>
- Gunaratne, C., Hatna, E., Epstein, J. M., & Garibay, I. (2023). Generating mixed patterns of residential segregation: An evolutionary approach. *Journal of Artificial Societies and Social Simulation*, 26(2), 7. <https://doi.org/10.18564/jasss.5081>
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *The American Economic Review*, 92(5), 1644–1655. <http://www.jstor.org/stable/3083270>
- Hu, Y., Bahamonde-Birke, F. J., & Ettema, D. (2025). Vehicle-to-grid, why not? An interview with battery electric vehicle users with various driving patterns in Utrecht, the Netherlands. *Transport Policy*, 164, 231–240. <https://doi.org/10.1016/j.tranpol.2025.02.009>
- Huang, B., Meijssen, A. G., Annema, J. A., & Lukso, Z. (2021). Are electric vehicle drivers willing to participate in vehicle-to-grid contracts? A context-dependent stated choice experiment. *Energy Policy*, 156, 112410. <https://doi.org/10.1016/j.enpol.2021.112410>
- Huber, J., Schaule, E., Jung, D., & Weinhardt, C. (2019). Quo vadis smart charging? a literature review and expert survey on technical potentials and user acceptance of smart charging systems. *World Electric Vehicle Journal*, 10(4), 85. <https://doi.org/10.3390/wevj10040085>

- Huda, M., Koji, T., & Aziz, M. (2020). Techno Economic Analysis of Vehicle to Grid (V2G) Integration as Distributed Energy Resources in Indonesia Power System. *Energies*, 13(5), 1162. <https://doi.org/10.3390/en13051162>
- International Energy Agency. (2023a). *Tracking clean energy progress 2023* (Licence: CC BY 4.0). International Energy Agency. Paris. <https://www.iea.org/reports/tracking-clean-energy-progress-2023>
- International Energy Agency. (2023b, October). *Electricity grids and secure energy transitions* (tech. rep.) (Accessed 2026-06-18). International Energy Agency. <https://www.iea.org/reports/electricity-grids-and-secure-energy-transitions>
- International Energy Agency. (2026, May). Vehicle-to-grid technology [Accessed 2026-06-22]. <https://www.iea.org/reports/vehicle-to-grid-technology>
- Jaxa-Rozen, M., & Kwakkel, J. (2018). Pynetlogo: Linking netlogo with python. *Journal of Artificial Societies and Social Simulation*, 21(2), 4. <https://doi.org/10.18564/jasss.3668>
- Jensen, T., & Chappin, E. (2017). Automating agent-based modeling: Data-driven generation and application of innovation diffusion models. *Environmental Modelling And Software*, 92, 261–268. <https://doi.org/https://doi.org/10.1016/j.envsoft.2017.02.018>
- Kempton, W., & Tomić, J. (2005). Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1), 268–279. <https://doi.org/10.1016/j.jpowsour.2004.12.025>
- KNMI. (2018). Daggegevens van het weer in nederland. <http://projects.knmi.nl/klimatologie/daggegevens/selectie.cgi>
- Kubli, M., Loock, M., & Wüstenhagen, R. (2018). The flexible prosumer: Measuring the willingness to co-create distributed flexibility. *Energy Policy*, 114, 540–548. <https://doi.org/10.1016/j.enpol.2017.12.044>
- Kwakkel, J. H. (2017). The exploratory modeling workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making. *Environmental Modelling & Software*, 96, 239–250. <https://doi.org/10.1016/j.envsoft.2017.06.054>
- Kwakkel, J. H., Haasnoot, M., & Walker, W. E. (2016). Coping with the wickedness of public policy problems: Approaches for decision making under deep uncertainty. *Journal of Water Resources Planning and Management*, 142(3), 01816001. [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000626](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000626)
- Lee, C.-Y., Jang, J.-W., & Lee, M.-K. (2020). Willingness to accept values for vehicle-to-grid service in South Korea. *Transportation Research Part D Transport and Environment*, 87, 102487. <https://doi.org/10.1016/j.trd.2020.102487>
- Lehtola, T. A. (2025). Vehicle-to-grid applications and battery cycle aging: A review. *Renewable and Sustainable Energy Reviews*, 208, 115013. <https://doi.org/https://doi.org/10.1016/j.rser.2024.115013>
- Lempert, R. J., Popper, S. W., & Bankes, S. C. (2003). *Shaping the next one hundred years: New methods for quantitative, long-term policy analysis* (tech. rep.). RAND Corporation. https://www.rand.org/pubs/monograph_reports/MR1626.html
- Li, K., Ge, J., Lomax, N., & Polhill, J. G. (2026). Calibrating a global trade agent-based model with an hpc-abc-smc framework. *Journal of Artificial Societies and Social Simulation*, 29(3), 5. <https://doi.org/10.18564/jasss.6072>
- Li, Y., Huang, Y., Liang, Y., Song, C., & Liao, S. (2024). Economic and carbon reduction potential assessment of vehicle-to-grid development in guangdong province. *Energy*, 302, 131742. <https://doi.org/10.1016/j.energy.2024.131742>
- Liu, F., Wei, Z., Lin, Y., Huang, X., Li, Y., Huang, Y., & Lim, M. K. (2025). Vehicle□To□Grid Technology Acceptance for Electric Vehicle Users: A Systematic Literature Review and Future Research Agenda. *International Journal of Consumer Studies*, 49(3). <https://doi.org/10.1111/ijcs.70065>
- MacKay, D. J. C. (1992). Bayesian interpolation. *Neural Computation*, 4(3), 415–447. <https://doi.org/10.1162/neco.1992.4.3.415>
- Mehdizadeh, M., Nordfjaern, T., & Klöckner, C. A. (2023). Estimating financial compensation and minimum guaranteed charge for vehicle-to-grid technology. *Energy Policy*, 180, 113649. <https://doi.org/10.1016/j.enpol.2023.113649>
- Meijssen, A. (2019). *Dutch electric vehicle drivers' preferences regarding vehicle-to-grid contracts* [Master's thesis]. Delft University of Technology.

- Meyer, J. T., Daam, C., & Gemassmer, J. (2022). Realise flexibility potential of ev fleets through grid-serving charging strategies. *Proceedings of the 6th E-Mobility Power System Integration Symposium*. https://reiner-lemoine-institut.de/wp-content/uploads/2023/04/Netz_eLOG_Delft2022_RLI_Print.pdf
- Milieu Centraal. (n.d.-a). Gemiddeld energieverbruik in nederland [Accessed 2026-06-22]. <https://www.milieucentraal.nl/energie-besparen/inzicht-in-je-energierekening/gemiddeld-energieverbruik/>
- Milieu Centraal. (n.d.-b). Opladen elektrische auto [Accessed 2026-06-22]. <https://www.milieucentraal.nl/duurzaam-vervoer/elektrische-auto/opladen-elektrische-auto/>
- Ministerie van Infrastructuur en Waterstaat. (2026). *Nationale routekaart bidirectioneel laden*. <https://elaad.nl/wp-content/uploads/2026/07/Nationale-Routekaart-Bidirectioneel-Laden.pdf>
- Mozafar, M. R., Amini, M. H., & Moradi, M. H. (2017). Innovative appraisal of smart grid operation considering large-scale integration of electric vehicles enabling V2G and G2V systems. *Electric Power Systems Research*, 154, 245–256. <https://doi.org/10.1016/j.epsr.2017.08.024>
- Nationale Agenda Laadinfrastructuur. (2025, December). Netbewust laden van start op 5.200 publieke laadpunten [Accessed: 2026-06-22]. <https://www.agendalaadinfrastructuur.nl/werkgroepen/werkgroep+smart+charging/nieuws+slim+laden+voor+iedereen/3177207.aspx>
- Netbeheer Nederland. (2023, June). *Het energiesysteem van de toekomst: De ii3050-scenario's. integrale energiesysteemverkenning 2030–2050* (tech. rep.) (Accessed 2026-06-18). Netbeheer Nederland. <https://www.netbeheernederland.nl/publicatie/rapport-ii3050-scenarios>
- Netbeheer Nederland. (2025, May). *Netbeheer nederland scenario's editie 2025* (tech. rep.) (Accessed 2026-06-18). Netbeheer Nederland. https://www.netbeheernederland.nl/sites/default/files/2025-05/netbeheer_nederland_scenarios_editie_2025_-_def_12-05.pdf
- Netbeheer Nederland. (2026). *Feiten en cijfers*. Retrieved August 24, 2026, from <https://www.netbeheernederland.nl/feiten-en-cijfers>
- Octopus Energy. (2026). *Power pack*. Retrieved August 19, 2026, from <https://octopus.energy/power-pack/>
- Ommering, N. (2024, July). *Unlocking solar potential: How inverse modelling can contribute to uncovering plausible explanations for residential solar pv adoption dynamics in the netherlands* [Master's thesis]. Delft University of Technology.
- Park Lee, E. H. (2019). *A socio-technical exploration of the car as power plant* [Doctoral dissertation, Delft University of Technology]. <https://doi.org/https://doi.org/10.4233/uuid:bcf14d54-74cc-4fe8-9285-466cee3936ab>
- Parsons, G. R., Hidrue, M. K., Kempton, W., & Gardner, M. P. (2014). Willingness to pay for vehicle-to-grid (V2G) electric vehicles and their contract terms. *Energy Economics*, 42, 313–324. <https://doi.org/10.1016/j.eneco.2013.12.018>
- Partners in Energie. (2026). *Capaciteitskaart elektriciteitsnet: Regionale netbeheerders – teruglevering*. Retrieved August 23, 2026, from <https://data.partnersinenergie.nl/capaciteitskaart/rnb/invoeding>
- Platt, D. (2020). A comparison of economic agent-based model calibration methods. *Journal of Economic Dynamics and Control*, 113, 103859. <https://doi.org/10.1016/j.jedc.2020.103859>
- Refa, N., Hammer, D., & Slobben, T. (2023, January). *Outlook laadprofielen personenauto's: Regulier en netbewust laden* (tech. rep.). ElaadNL. Arnhem. https://elaad.nl/wp-content/uploads/downloads/Rapport_Outlook_Laadprofielen.pdf
- Renzini, F., Debernardi, C., Bianchi, F., Cremonini, M., & Squazzoni, F. (2023). The new frontiers of social simulation in the data science era: An introduction to the proceedings. In F. Squazzoni (Ed.), *Advances in social simulation* (pp. 1–10). Springer Nature Switzerland.
- Rijksdienst voor Ondernemend Nederland. (2025, May). *Tendrapport zware bedrijfsvoertuigen: Editie 2025* (tech. rep.) (Accessed 2026-06-22). Rijksdienst voor Ondernemend Nederland. <https://www.rvo.nl/sites/default/files/2025-05/Tendrapport%20Zware%20Bedrijfsvoertuigen%20Editie%202025.pdf>
- Rijksdienst voor Ondernemend Nederland. (2026a, April). *Implementatie van grootschalige batterijen in nederland* (tech. rep.) (Accessed 2026-06-22). Rijksdienst voor Ondernemend Nederland. <https://www.rvo.nl/sites/default/files/2026-04/Implementatie%20van%20grootschalige%20batterijen%20in%20Nederland.pdf>

- Rijksdienst voor Ondernemend Nederland. (2026b, April). Stand van zaken elektrisch vervoer en laadpunten [Vehicle fleet data for March 2026. Accessed 2026-06-22]. <https://www.rvo.nl/onderwerpen/elektrisch-vervoer/stand-van-zaken>
- Rijksoverheid. (n.d.). Overheid stimuleert milieuvriendelijker rijden [Accessed 2026-06-22]. <https://www.rijksoverheid.nl/onderwerpen/auto/overheid-stimuleert-milieuvriendelijker-rijden>
- Sagaria, S., Van Der Kam, M., & Boström, T. (2024a). Conceptualization of a vehicle-to-grid assisted nation-wide renewable energy system – A case study with Spain. *Energy Conversion and Management X*, 22, 100545. <https://doi.org/10.1016/j.ecmx.2024.100545>
- Sagaria, S., Van Der Kam, M., & Boström, T. (2024b). The influence of socio-technical variables on vehicle-to-grid technology. *Energy*, 305, 132299. <https://doi.org/10.1016/j.energy.2024.132299>
- Schubotz, L. (2025). Parallel sampler [Accessed: 2026-03-24]. <https://gitlab.tudelft.nl/lshubotz/Parallel-Sampler/>
- Schubotz, L., Chappin, E., & Scholz, G. (2025). Structural sculpting: Making inverse modelling generate and deal with variable structures. *Journal of Artificial Societies and Social Simulation*, 28(4), 10. <https://doi.org/10.18564/jasss.5830>
- Schubotz, L., Chappin, E., Scholz, G., & de Vries, G. (2026). Participatory inverse modelling for behavioural insights into Dutch EV charging dynamics. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-56059-7>
- scikit-learn API. (2026). Bayesianridge [Accessed 2026-06-18]. https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.BayesianRidge.html
- Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. In F. Pereira, C. Burges, L. Bottou, & K. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 25). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf
- TenneT. (n.d.). Stikstof [Accessed 2026-06-22]. <https://www.tennet.eu/nl/over-tennet/organisatie-en-strategie/onze-standpunten/stikstof>
- V2G Hub. (2026). *Insights*. Retrieved August 21, 2026, from <https://www.v2g-hub.com/insights/>
- Van Der Kam, M., Peters, A., Van Sark, W., & Alkemade, F. (2019). Agent-Based Modelling of Charging Behaviour of Electric Vehicle Drivers. *Journal of Artificial Societies and Social Simulation*, 22(4). <https://doi.org/10.18564/jasss.4133>
- van der Kam, M., & van Sark, W. (2015). Smart charging of electric vehicles with photovoltaic power and vehicle-to-grid technology in a microgrid; a case study. *Applied Energy*, 152, 20–30. <https://doi.org/10.1016/j.apenergy.2015.04.092>
- Van Heuveln, K., Ghotge, R., Annema, J. A., Van Bergen, E., Van Wee, B., & Pesch, U. (2021). Factors influencing consumer acceptance of vehicle-to-grid by electric vehicle drivers in the Netherlands. *Travel Behaviour and Society*, 24, 34–45. <https://doi.org/10.1016/j.tbs.2020.12.008>
- Vattenfall. (2026). *Vattenfall start v2g-pilot: Elektrische auto als buffer voor het stroomnet*. Retrieved August 20, 2026, from <https://group.vattenfall.com/nl/newsroom/persbericht/2026/vattenfall-start-v2g-pilot-elektrische-auto-als-buffer-voor-het-stroomnet>
- Vey, M., Ruhnau, O., & Namockel, N. (2025). Green Charged Decision-Making: How Two-Part Remuneration, Contract Flexibility, and Environmental Nudging Drive Vehicle-to-Grid Participation. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.5374868>
- Voulis, N., Warnier, M., & Brazier, F. M. T. (2017). Impact of service sector loads on renewable resource integration. *Applied Energy*, 205, 1311–1326. <https://doi.org/10.1016/j.apenergy.2017.07.134>
- Vu, T. M., Buckley, C., Duro, J. A., Brennan, A., Epstein, J. M., & Purhouse, R. (2023). Can social norms explain long-term trends in alcohol use? insights from inverse generative social science. *Journal of Artificial Societies and Social Simulation*, 26(2), 4. <https://doi.org/10.18564/jasss.5077>
- Wang, W., Li, K., Liu, Y., Lian, J., & Hong, S. (2022). A system dynamics model analysis for policy impacts on green agriculture development: A case of the Sichuan Tibetan Area. *Journal of Cleaner Production*, 371, 133562. <https://doi.org/10.1016/j.jclepro.2022.133562>
- Wilensky, U., & Rand, W. (2015). *An introduction to agent-based modeling: Modeling natural, social, and engineered complex systems with netlogo*. MIT press.
- Will, C., & Schuller, A. (2016). Understanding user acceptance factors of electric vehicle smart charging. *Transportation Research Part C Emerging Technologies*, 71, 198–214. <https://doi.org/10.1016/j.trc.2016.07.006>

- Yaman, F., Yakhno, V. G., & Potthast, R. (2013). A survey on inverse problems for applied sciences. *Mathematical Problems in Engineering*, 2013, 976837. <https://doi.org/10.1155/2013/976837>
- Yun, S., Woo, J., & Kwak, K. (2025). Unlocking peak shaving: How EV driver heterogeneity shapes V2G potential. *Energy*, 329, 136773. <https://doi.org/10.1016/j.energy.2025.136773>
- Zhang, C., Kitamura, H., & Goto, M. (2024). Exploring V2G Potential in Tokyo: The Impact of User Behavior Through Multi-Agent Simulation. *IEEE Access*, 12, 118981–119002. <https://doi.org/10.1109/access.2024.3449448>
- Zhang, C., Kitamura, H., & Goto, M. (2026). Behavioral insights into vehicle-to-grid (V2G) adoption: Beyond economic perspectives. *Renewable and Sustainable Energy Reviews*, 228, 116603. <https://doi.org/10.1016/j.rser.2025.116603>
- Zonneveld, J. (2019). *Increasing participation in v2g through contract elements* [Master's thesis]. Delft University of Technology.



Evidence on V2G participation preferences and contract design

This appendix provides the study-specific evidence underlying the discussion of V2G user preferences and contract design in Chapter 2. It summarises the contract attributes, experimental levels, main findings, and reported preference heterogeneity in selected stated-choice and related participation studies.

A.1. Behavioural foundations of participation

A systematic review of 87 studies by Liu et al. (2025) shows that V2G adoption is shaped by three broad categories of determinants: product characteristics, such as battery degradation and charging flexibility; individual perceptions, including range anxiety and perceived risk; and economic incentives. The review also identifies a broad range of theoretical approaches, with demand-response models being the most common. These models generally represent participation as a response to prices or financial incentives and are useful for modelling demand-side flexibility and price responsiveness, but they typically do not explicitly represent user preferences, perceptions, or social influences. Despite growing recognition of these behavioural factors, the literature therefore remains strongly focused on technical system design and optimisation, with comparatively limited attention to how users evaluate V2G participation and how their preferences, perceptions, and experiences influence participation decisions.

The second most frequently applied framework identified in the review is the Technology Acceptance Model (TAM), which explains technology adoption through perceived usefulness, perceived ease of use, attitudes, and behavioural intentions (Davis et al., 1989). In an extended TAM study of EV charging scheduling, Wang et al. (2022) confirm the importance of these mechanisms and show that trust, social value, social influence, and perceived risk affect acceptance through different pathways. Trust and social influences can strengthen perceived usefulness, ease of use, and attitudes, whereas perceived risk directly reduces behavioural intention.

Beyond TAM, several studies draw on the Theory of Planned Behaviour (TPB) and related social-psychological approaches to explain participation in smart charging and V2G (Ghotge et al., 2022; Van Heuveln et al., 2021; Will & Schuller, 2016). Across these studies, acceptance is strengthened by perceived system benefits, particularly contributions to grid stability and renewable-energy integration, as well as by transparent operation, user insight, and control over charging and discharging. Evidence from real-world pilots further shows that users value clear information about battery state of charge and degradation, minimum battery settings, and the possibility to opt out (Ghotge et al., 2022). Conversely, restrictions on vehicle use, range anxiety, perceived loss of freedom, inconvenience, and concerns about battery degradation are repeatedly identified as barriers to acceptance (Ghotge et al., 2022; Van Heuveln et al., 2021; Will & Schuller, 2016). Experience with V2G can improve understanding and help users develop more concrete preferences regarding acceptable participation conditions (Ghotge et al., 2022). The effect of financial compensation is less consistent across these studies and is discussed further in the next section.

Other behavioural and economic approaches reinforce these findings. Models combining the Theory of Reasoned Action (TRA), TAM, and econometric analysis identify minimum guaranteed battery levels, compensation, trust, perceived usefulness, EV experience, and demographic characteristics as relevant to participation (Mehdizadeh et al., 2023). Contingent-valuation research shows that compensation alone may be insufficient: more than half of the respondents studied by Lee et al. (2020) rejected participation because of the obligation to connect their vehicle when requested.

Random-utility approaches instead represent participation as a trade-off between contract attributes, with users selecting the alternative that provides the highest utility (Huang et al., 2021; Zonneveld, 2019). This utility-based decision logic is also adopted in this thesis and embedded within the agent-based model, where heterogeneous contract preferences determine participation and can be influenced by evolving behavioural factors. Recent research using the Capability, Opportunity, Motivation–Behaviour (COM-B) model indicates that charging opportunities and motivational factors such as environmental benefits, grid stability, financial value, and trust may be more influential than users’ capabilities (De Graaf et al., 2026).

Importantly, positive attitudes towards V2G do not necessarily translate into actual participation. Zhang et al. (2026) distinguish psychological acceptance from concrete engagement and show that users may support V2G in principle while refraining from participation because of perceived risks, uncertainty about system performance, or mobility constraints. More broadly, EV-adoption research shows that participation-related preferences also vary with socio-demographic characteristics, sustainability values, mobility practices, and previous EV experience (Chen et al., 2020). Potential V2G revenues may increase the attractiveness of EVs, but economic incentives remain embedded in this wider behavioural context.

Taken together, the literature supports modelling V2G participation as a heterogeneous and potentially dynamic decision process. Participation is shaped by the interaction between contract conditions, mobility needs, perceived usefulness and risks, user control, experience, and social influence. These mechanisms provide the behavioural foundation for this thesis. The next section examines stated-choice evidence in more detail to determine how users trade off the specific contract attributes associated with V2G participation.

A.2. Contract attributes and experimental levels

Stated-choice studies examine V2G preferences by asking respondents to choose between hypothetical contracts that differ in one or more attributes. The selected studies operationalise these contracts in different ways, but commonly include financial remuneration, minimum State of Charge (SOC) or driving-range guarantees, required plug-in duration or frequency, battery degradation, contract duration, and forms of user control.

Table A.1 provides an overview of the attributes and levels included in each study, together with the country in which the research was conducted. Monetary values, time units, and technical restrictions are reported in the form used by the original studies and have not been converted to a common unit. Empty cells indicate that the corresponding attribute was not included in the reported experiment. Because the studies differ in samples, national contexts, experimental designs, and attribute definitions, their numerical results are not directly comparable. Statements about relative importance should therefore be interpreted within the specific choice setting of each study.

Table A.1: Contract attributes and experimental levels used in selected studies of EV charging and V2G participation

	Vey et al. (2025), Germany	Parsons et al. (2014), US	Geske and Schumann (2018), Germany	Zonneveld (2019), Netherlands	Kubli et al. (2018), Swiss	Meijssen and Huang et al. (2021), Netherlands	Bakhuys et al. (2025), Netherlands	Yun et al. (2025), Korea
Renumeration fixed	0, 25, 50 cent/day	500, 1000, 2000, 3000, 4000, 5000 USD/year	10-60 Euro/month		50, 70, 90, 110 CHF/month	20, 60, 100 Euro/month	40, 90, 140 Euro/month	25, 42, 58 USD/month
Refumeration variable	5, 10, 15 cent/kWh			20, 60, 100 cent/hour		No, 15, 30 cent/hour		
Contract flexibility	Day or month			1, 12, 24 months	0, 12, 24, 48 months	6, 12, 24 months		
Days plugged in			Between 3 and 5/week			4, 7, 10 times month maximum drained		4, 8, 12, 16 workdays
Min-plug in duration	5, 10, 15 hours/day	5, 10, 15, 20 hours/day	Between 0 and 14 hours/day	0, 25, 30 hours/week		5, 10, 15 hours/day	2, 5, 8 hours per day	
Min battery level	35, 50, 75%				40, 60, 80, 100%			Same as arrival, 50, 80%
Minimum range		25, 75, 125, 175 miles	Between 10 to 50 km	10, 50, 90 km			25, 75, 125 km	
Price relative to preferred GV		1000, 2000, 3000, 4000, 8000, 16000 dollar higher	Between 0 and 7000 euro price reduction					
Battery degradation				1, 4, 7 cycles per session	0, 1, 3, unlimited cycles day	1, 4, 7 cycles daily	1, 2, 3% yearly	
Distance to charger								0, 1, 2 km
Connection time								8-12, 2-6, 6-10, 10-6
App/trip control			Yes/no					

A.3. Study-specific findings and preference heterogeneity

The use of similar contract attributes does not necessarily produce identical findings across studies. Their effects depend on the levels examined, the other attributes included in the choice set, the charging context, and differences between respondents. Table A.2 therefore summarises the principal contract findings of each study together with reported heterogeneity and relevant qualifications. Terms such as “most important” or “primary driver” refer to the results within the respective experimental design and should not be interpreted as directly comparable effect sizes across studies.

Table A.2: Main findings from selected stated-choice studies on V2G contract preferences and heterogeneity

Study	Main findings on contract attributes	Heterogeneity and qualifications
Parsons et al. (2014)	Financial compensation is the most influential attribute, while mobility guarantees are also a major determinant of participation and rank second in importance.	Younger respondents, men, respondents planning to purchase a hybrid vehicle, and those expecting access to a home charging outlet are more likely to belong to the EV/V2G-oriented class. Income does not significantly affect class membership. Respondents reporting major or minor pro-environmental lifestyle changes are approximately three and two times, respectively, as likely to belong to this class as respondents reporting no such changes.
Kubli et al. (2018)	Compensation is a primary driver of contract acceptance. Flexibility is represented by a combined attribute covering minimum State of Charge (SOC) and the number of discharge cycles. Its overall importance is moderate, but acceptance declines sharply when guaranteed SOC falls below approximately 60%. Shorter contracts are preferred.	Contract duration is examined within a broader flexibility proposition that also includes PV, storage, and heat-pump integration. This wider context may increase the perceived commitment associated with longer contracts.
Zonneveld (2019)	Compensation strongly affects contract preferences. Plug-in duration is not statistically significant, possibly because the experiment uses relatively short connection durations. Contract duration ranks as the third most important attribute, with longer contracts preferred. The number of discharge cycles has an overall negative effect on utility, but the estimated non-linear pattern is inconsistent: utility falls between one and four cycles and unexpectedly rises at seven cycles.	The positive effect found for battery degradation is difficult to interpret and may reflect respondents' difficulty understanding degradation expressed in discharge cycles rather than a genuine preference. Longer contracts may have been perceived as offering stability and reducing the need for repeated renegotiation.
Geske and Schumann (2018)	Minimum guaranteed range is the most important attribute and is more influential than compensation. Some user groups also attach substantial importance to maintaining flexibility in their charging behaviour.	Range anxiety distinguishes more sceptical user groups that are less willing to participate. Income and environmental attitudes have limited explanatory power and do not consistently determine segment membership.
Meijssen (2019); Huang et al. (2021)	Under normal recharging, guaranteed minimum battery level has the highest relative importance, followed by plug-in time; under hypothetical fast recharging, plug-in time becomes most important as the importance of the battery guarantee decreases substantially. Fixed remuneration positively affects participation, while longer plug-in times and more discharging cycles reduce utility. Variable extra remuneration and contract duration are statistically insignificant.	The results show that contract preferences depend on the surrounding charging technology: fast recharging substantially reduces the importance attached to the guaranteed battery level.
Bakhuis et al. (2025)	Guaranteed range dominates preferences for the two largest user groups, representing 43% and 29% of respondents, and accounts for approximately 60% and 45% of attribute importance in these groups. A flexibility-oriented segment, representing approximately 18% of respondents, assigns more than half of attribute importance to minimum plug-in time. Only a smaller group of approximately 10% places strong emphasis on compensation. Battery degradation is significant but is not the primary concern of any group.	The results reveal distinct mobility-, compensation-, and flexibility-oriented preference segments. Range anxiety helps distinguish more sceptical users, whereas income and environmental attitudes do not consistently explain class membership.

Table A.2 continued

Study	Main findings on contract attributes	Heterogeneity and qualifications
Vey et al. (2025)	Reducing guaranteed SOC from 70% to 50% does not conclusively affect participation, whereas a reduction to 35% is clearly unacceptable. Moderate changes in per-kWh remuneration have little effect, whereas very low payment levels substantially reduce utility, suggesting a possible threshold between 5 and 10 cents/kWh. Daily contract flexibility is strongly preferred over monthly commitment. Longer minimum plug-in durations reduce participation, particularly at 15 hours.	Responses to remuneration and contract restrictions vary with income. Higher-income respondents are more responsive to remuneration changes and tend to prefer flexible contracts and higher guaranteed battery levels, whereas lower-income respondents require greater compensation as plug-in requirements increase. Broader evidence links income to risk attitudes, although income effects are not consistent across all V2G studies (Holt & Laury, 2002).
Yun et al. (2025)	The timing of plug-in events is an important determinant of contract preference. Compensation class dependent.	Preferences for plug-in timing differ substantially across users. Approximately one third of respondents assign very high importance to timing, at around 50% relative importance, while another third assign moderate importance, at around 25%.

B

Experimental design and computational settings

This appendix provides the detailed settings underlying the methodology described in Chapter 3. It compares the exploratory and final parameter spaces, lists the model settings held constant during the analysed simulations, and reports the computational configuration used in the final learning phase.

B.1. Exploratory and final parameter spaces

Table B.1 reports the discrete parameter values used during the exploratory analysis and the final learning phase. The exploratory space was deliberately broad and contained 334,353,766,809,600 possible combinations. Based on the exploratory results, the final space was reduced to 75,937,500 combinations. Parameters reported as fixed were varied during the exploratory phase but held constant during the final learning phase.

Table B.1: Exploratory and final parameter spaces.

Category	Parameter	Exploratory values	Final values	Basis
Contract	contract-FC	$n = 10$: 30, 50, 60, 75, 90, 120, 130, 150, 170, 200	$n = 6$: 30, 60, 90, 120, 150	Bakhuis et al.; exploratory screening
Contract	contract-GR	$n = 16$: 10, 15, 40, 45, 60, 75, 105, 110, 125, 135, 150, 160, 165, 170, 185, 200	$n = 6$: 15, 45, 75, 105, 135	Bakhuis et al.; exploratory screening
Contract	contract-MP	$n = 12$: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12	$n = 5$: 1, 3, 5, 7, 9	Bakhuis et al.; exploratory screening
Contract	contract-BD	$n = 6$: 1, 1.5, 2, 2.5, 3, 4	$n = 3$: 1, 2, 3, 4	Bakhuis et al.; exploratory screening
Charging infrastructure	private-chargers-slider	$n = 10$: 0.5, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.4, 1.5, 1.7	$n = 5$: 0.8, 0.9, 1.0, 1.1, 1.2	Centred around reference value
Charging infrastructure	public-chargers-slider	$n = 9$: 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.5, 1.8	$n = 5$: 0.8, 0.9, 1.0, 1.1, 1.2	Centred around reference value
Charging infrastructure	office-chargers-slider	$n = 8$: 0.8, 0.9, 1.0, 1.1, 1.2, 1.7, 1.9, 2.0	$n = 5$: 0.8, 0.9, 1.0, 1.1, 1.2	Centred around reference value
Charging infrastructure	commercial-chargers-slider	$n = 9$: 0.5, 0.8, 0.9, 1.0, 1.1, 1.2, 1.5, 1.6, 1.7	$n = 5$: 0.8, 0.9, 1.0, 1.1, 1.2	Centred around reference value
Renewable energy	renewable-coverage	$n = 9$: 40, 50, 60, 66, 83, 91, 92, 93, 98	$n = 3$: 40, 50, 60	Bakhuis et al.; exploratory screening
Renewable energy	solar-vs-wind	$n = 9$: 40, 50, 60, 65, 77, 80, 88, 95, 100	$n = 3$: 40, 50, 60	Bakhuis et al.; exploratory screening
Charging policy	home-charging-policy	$n = 3$: normal, net-aware, smart	$n = 3$: normal, net-aware, smart	Policy alternatives
Charging policy	work-charging-policy	$n = 3$: normal, net-aware, smart	$n = 3$: normal, net-aware, smart	Policy alternatives
Charging policy	public-charging-policy	$n = 3$: normal, net-aware, smart	$n = 3$: normal, net-aware, smart	Policy alternatives
V2G control	charge-reduce-start-kW	$n = 8$: 325, 400, 450, 500, 525, 550, 575, 600	$n = 1$: fixed at 450	Exploratory loss and plausibility checks
V2G control	charge-stop-kW	$n = 8$: 550, 600, 650, 675, 700, 725, 750, 775	$n = 1$: fixed at 650	Exploratory loss and plausibility checks
V2G control	discharge-start-kW	$n = 4$: 500, 550, 600, 625	$n = 1$: fixed at 550	Exploratory loss and plausibility checks
V2G control	full-discharge-kW	$n = 8$: 600, 650, 675, 700, 750, 775, 800, 825	$n = 1$: fixed at 750	Exploratory loss and plausibility checks
Total parameter-space size		334,353,766,809,600	75,937,500	

Notes: n denotes the number of discrete values. Unless otherwise indicated, parameters and alternatives were introduced as part of the model extension developed in this study. “Bakhuis et al.” identifies values derived from the underlying V2G model and stated-choice design.

B.2. Fixed parameter settings

Table B.2 lists the parameters held constant throughout the analysed simulation runs. Unless stated otherwise, these settings were inherited from the model developed by Bakhuis (2026).

Table B.2: Fixed model settings used in the analysed simulations.


Category	Parameter	Fixed value
Core scenario	charge-mode	non-V2G
Core scenario	ev-types	fev - all the same
		battery capacity
Core scenario	charging-scenario	per-location
Vehicle	battery-capacity	100
Vehicle	charge-speed	10
Market and pricing	imbalance-threshold-kW-hour	10
Market and pricing	pos-price-slope	0.003
Market and pricing	neg-price-slope	0.01
Aggregator and contract	aggregator?	true
Aggregator and contract	contract-re-evaluation?	true
Aggregator and contract	fc-step	5
Aggregator and contract	fc-min	10
Aggregator and contract	fc-max	200
Aggregator and contract	update-GR-and-MP?	false
Aggregator and contract	gr-step	5
Aggregator and contract	gr-min	10
Aggregator and contract	gr-max	130
Aggregator and contract	mp-step	1
Aggregator and contract	mp-min	2
Aggregator and contract	mp-max	12
Psychological mechanisms	psychology?	true
Psychological mechanisms	use-ESI?	true
Psychological mechanisms	use-RA?	true
Psychological mechanisms	use-PIC?	true
Environmental self-identity	f-esi-plus	1×10^{-4}
Environmental self-identity	f-esi-min	5×10^{-4}
Environmental self-identity	esi-mean	0.0
Environmental self-identity	esi-sd	0.333
Range anxiety	f-ra-plus	0.05
Range anxiety	f-ra-min	1×10^{-4}
Range anxiety	ra-mean	0.5
Range anxiety	ra-sd	0.25
Plug-in time concern	f-pic-plus	0.002
Plug-in time concern	f-pic-minus	0.0015
Plug-in time concern	pic-mean	0.5
Plug-in time concern	pic-sd	0.15
Social influence	peer-influence?	false
Social influence	num-peers	0
Social influence	max-peer-threshold-shift	0.10

B.3. Computational environment and final learning settings

The final learning experiment series used a common computational configuration. Table B.3 reports the software environment, parallel execution, simulation, randomisation, and sampler settings.

Table B.3: Computational environment and common settings used during the final learning phase.

Setting	Implementation
Computing environment	DelftBlue high-performance computing cluster at Delft University of Technology.
Job scheduler	Slurm.
Allocation per job	One node, one task, and four CPUs per task.
Parallel worker processes	Four independent Python worker processes, each operating an independent headless NetLogo instance.
Maximum wall time	Up to 23 hours per submitted job.
Python version	Python 3.10.12.
Java version	OpenJDK 17.0.8.1.
NetLogo version	NetLogo 6.4.0.
Execution mode	NetLogo was operated in headless mode through the Python wrapper.
Simulation horizon	96,768 ticks per individual simulation run.
Surrogate model	Bayesian Ridge regression.
Final sampling method	Exploitative sampling: candidate configurations were ranked by their predicted loss, and configurations with the lowest predicted loss were selected.
Candidate evaluation	Four parameter configurations were evaluated before each surrogate-model update, corresponding to the number of available worker processes.
Repetitions	Ten stochastic repetitions per parameter configuration.
Simulations per complete Bayesian update	Four parameter configurations multiplied by ten repetitions, resulting in 40 individual NetLogo runs.
Task chunk size	One simulation task per worker assignment to support load balancing between runs with different computational durations.
Wrapper-level randomness	A separate Python-wrapper seed was configured for each learning stream
NetLogo randomness	Each repetition received an explicit NetLogo random seed. Within an experiment series, the same repetition-seed set was reused across parameter configurations, and the realised seed was stored with each simulation output.
Preloading	Preloading was enabled. V2G 1 and V2G 2 used separate cumulative training histories.
Training-history locations	Separate history directories were maintained for V2G_1 and V2G_2 on the DelftBlue scratch filesystem.
Result storage	Outputs returned for a completed Bayesian batch were combined and stored in a separate pickle file containing the parameter settings, repetition identifiers, random seeds, model reporters, and calculated losses.
Timeout handling	Completed and readable batch files from timed-out jobs were retained when they contained all expected simulation outputs and valid loss values. Incomplete or corrupted batch files were excluded.
Initial random design size	n_{10} randomly sampled parameter configurations.
Candidate pool size	n_{100} previously unevaluated configurations generated and scored per Bayesian update.



Model

This appendix provides additional documentation of the agent-based model underlying the methodology described in Chapter 4. It summarises the inherited model foundations, defines the principal agent and environment variables, and describes the charging, behavioural-feedback, and V2G participation procedures in greater detail.

Tables C.1–C.3 define the principal state variables of the EV-driver agents, grid cells, and simulation environment, respectively.

C.1. Inherited model foundations

This section provides additional background on the two successive agent-based models underlying the model used in this study. It first summarises the original EV charging model developed by Van Der Kam et al. (2019) and then describes the V2G-oriented extension developed by Bakhuis (2026). The adaptations introduced in the present study are discussed in Section 4.3.

C.1.1. Original model by Van der Kam et al.

The original model by Van der Kam et al. is not used directly in this study, but forms the foundation of the V2G model developed by Bakhuis. It is an agent-based model of EV charging behaviour in a stylised urban electricity system, developed to compare interventions for more sustainable charging behaviour, including financial incentives, automated smart charging, information provision, and social charging.

Its main contribution for the later V2G model is the basic transport and electricity structure. EV drivers follow diary-based mobility patterns, interact with charging infrastructure at residential, office, and commercial locations, and charge under time-varying electricity demand and renewable generation conditions. The model also already includes psychological drivers of charging behaviour, especially environmental self-identity and range anxiety, which influence how sustainably agents charge under information and feedback interventions.

However, the original model does not include bidirectional charging, V2G contract adoption, guaranteed range constraints, minimum plug-in requirements, battery degradation compensation, or an aggregator that coordinates charging and discharging. These components were introduced in the Bakhuis extension, which is the direct starting point for the present study.

C.1.2. V2G extension by Bakhuis et al.

The direct starting point for the present study is the V2G-oriented extension developed by Bakhuis. This extension builds on the Van der Kam et al. model by adding bidirectional charging, V2G contract evaluation, endogenous participation, driver heterogeneity, and an aggregator agent. It therefore shifts the model from EV charging behaviour and smart-charging interventions towards a sociotechnical representation of V2G adoption and system operation.

A key contribution of the Bakhuis extension is the representation of V2G contract choice. Drivers evaluate contracts based on financial compensation, guaranteed range, minimum plug-in time, and battery degradation. These preferences are empirically grounded in a stated-choice experiment and represented through latent preference classes. Participation is therefore not imposed as a fixed input, but emerges from heterogeneous driver preferences, participation thresholds, and contract evaluation.

The Bakhuis extension also repositions the psychological mechanisms from the original charging model. Range anxiety and environmental self-identity were already present in Van der Kam et al., but are used in the V2G extension to influence contract preferences and participation rather than directly determining charging mode. In addition, Bakhuis et al. introduce plug-in concern to represent the perceived inconvenience of meeting minimum plug-in requirements and peer-pressure.

Finally, the Bakhuis model adds an aggregator that offers V2G contracts and coordinates the charging and discharging of participating vehicles. Participating vehicles are operated within contractual constraints, especially the guaranteed range, while non-participants continue to charge independently. This creates the V2G participation and control structure that forms the direct basis for the present congestion-focused adaptation.

C.2. Agent variables

Table C.1: State variables of the EV-driver agent

Name	Description	Domain	Dynamic
Location	Current grid cell occupied by the EV driver	Patch coordinates	Y
Battery capacity	Usable battery capacity of the assigned vehicle	kWh	N
Home	Residential patch assigned as the driver's home	Patch	N
Home-charge?	Whether a private charging point is available at the driver's home	{true, false}	N
Next trip	Scheduled departure time of the next trip	Simulation time	Y
Battery energy	Energy currently stored in the vehicle battery	[0, battery capacity] kWh	Y
Minimum battery energy	Minimum battery energy retained for mobility or the guaranteed-range requirement	[0, 0.95 battery capacity] kWh	Y
Charging?	Whether the vehicle is charging during the current tick	{true, false}	Y
Discharging?	Whether the vehicle is discharging during the current tick	{true, false}	Y
Charging mode	Current participation mode: non-V2G or V2G	{1,2}	Y
Plug-in hours	Accumulated connected time during the daily 07:00–19:00 assessment period	Hours	Y
Environmental self-identity	Driver-specific environmental self-identity score	[−1, 1]	Y
Range anxiety	Driver-specific concern about insufficient battery energy	[0, 1]	Y
Plug-in concern	Driver-specific concern about long or restrictive plug-in requirements	[0, 1]	Y
Latent preference class	Empirically derived preference class determining initial contract sensitivities	{1, 2, 3, 4}	N
Guaranteed range	Driving range guaranteed under the evaluated V2G contract	km	N
Minimum plug-in time	Minimum daily plug-in duration required by the evaluated V2G contract	Hours	N
Battery degradation	Additional annual battery degradation associated with the evaluated contract	Percentage points	N

Name	Description	Domain	Dynamic
β_{FC}	Driver-specific sensitivity to financial compensation	Non-negative real	N
β_{GR}	Driver-specific sensitivity to guaranteed range	Positive real; class-constrained	Y
β_{MP}	Driver-specific sensitivity to minimum plug-in time	Negative real; class-constrained	Y
β_{BD}	Driver-specific sensitivity to battery degradation	Non-positive real	N
Utility score	Utility assigned to the currently evaluated V2G contract	Real	Y
Participation threshold	Minimum utility required for V2G participation	Real; class-constrained	Y
Social peers	Other EV drivers included in the driver's social network	Agentset of EV drivers	N
Peer satisfaction	Mean participation margin of adopting peers	Real	Y
Peer-threshold shift	Adjustment of the participation threshold caused by peer influence	$[-m, m]$, where m is the maximum peer shift	Y
Smart-charging target	Target battery level used by the smart-charging policy	kWh	N
Requested set-point	Requested charging or discharging energy during the current tick; negative values indicate discharge	kWh per tick	Y
Applied setpoint	Energy actually transferred after charger and battery constraints are applied	kWh per tick	Y
Grid-supportive charging	Whether the energy transfer during the current tick supported the system energy balance	{true, false}	Y

Table C.2: State variables of grid cells

Name	Description	Domain	Dynamic
Location	Coordinates of the grid cell	Coordinates	N
Function	Land-use or charging function assigned to the grid cell	Residential, office, commercial, house, or location-specific charging-station category	N
Charging point?	Whether the grid cell contains a charging point	{true, false}	N
Occupied?	Whether a semi-public charging point is currently occupied	{true, false}	Y
Reserved?	Whether a public charging point is reserved by an approaching EV driver	{true, false}	Y
Charging demand	Energy supplied to connected EVs at the grid cell during the current tick	kWh per tick	Y
Discharging demand	Energy supplied by connected V2G vehicles at the grid cell during the current tick	kWh per tick	Y

Table C.3: State variables of the simulation environment

Name	Description	Domain	Dynamic
Simulation time	Number of five-minute simulation steps elapsed	Ticks	Y
Solar capacity	Installed photovoltaic capacity implied by renewable coverage and the solar–wind allocation	MWp	Y
Photovoltaic production	Solar energy produced during the current tick	kWh per tick	Y
Wind capacity	Installed wind-generation capacity implied by renewable coverage and the solar–wind allocation	MW	Y
Wind production	Wind energy produced during the current tick	kWh per tick	Y
Residential electricity demand	Residential-sector electricity consumption during the current tick	kWh per tick	Y
Service-sector electricity demand	Service-sector electricity consumption during the current tick	kWh per tick	Y
Baseline energy balance	Renewable production minus residential and service-sector demand before EV charging and discharging	kWh per tick	Y
Energy balance	Remaining system energy balance after EV charging and discharging	kWh per tick	Y
EV charging demand	Aggregate EV charging demand during the current tick	kW	Y
EV discharging supply	Aggregate electricity supplied by V2G vehicles during the current tick	kW	Y
Non-V2G charging demand	Aggregate charging demand of non-V2G vehicles during the current tick	kW	Y
Charged energy by location	Energy charged at home, work, public, and commercial charging locations during the current tick	kWh per tick	Y
Discharged energy by location	Energy discharged at home, work, public, and commercial charging locations during the current tick	kWh per tick	Y
Driven electrically	Distance travelled with sufficient battery energy during the current tick	Patch-distance units	Y
Driven with low battery	Distance travelled with a battery state at or below 10% during the current tick	Patch-distance units	Y
Driven with empty battery	Distance travelled after battery energy has reached zero during the current tick	Patch-distance units	Y
Number of EV drivers	Number of EV-driver agents included in the simulation	Positive integer	N
Charging speed	Maximum charging power available to an EV	kW	N
Discharging speed	Maximum V2G discharging power; set equal to charging speed in the analysed model	kW	N
Private charging stations	Number of private residential charging points	Non-negative integer	N
Public charging stations	Number of public residential charging points	Non-negative integer	N
Office charging stations	Number of semi-public office charging points	Non-negative integer	N
Commercial charging stations	Number of semi-public commercial charging points	Non-negative integer	N

Name	Description	Domain	Dynamic
Imbalance threshold	Energy-balance threshold used by the aggregator to initiate control actions	kWh per tick	N
Loss threshold	Net-load threshold above which congestion loss is accumulated	kW	N

C.3. Detailed charging behaviour

This section provides the detailed charging and discharging procedures underlying the summary in Section 4.5. It specifies charging eligibility, the order in which charging decisions are executed, and the distinction between settings explored as experimental parameters and settings held fixed throughout the final analysis.

C.3.1. Charging access and execution order

A vehicle can use a charging point only after reaching its current destination. Access additionally depends on the presence and availability of a suitable charger. The reservation procedure is applied only after the vehicle has satisfied the corresponding charging conditions, preventing vehicles that are not eligible to charge from unnecessarily reserving a charging point.

Charging decisions are evaluated in five-minute ticks. Non-V2G charging is applied according to the policy associated with the vehicle's charging location. V2G control is subsequently applied to connected participating vehicles using the updated system state. Each V2G vehicle has one applied setpoint per tick, preventing simultaneous charging and discharging. All charging and discharging quantities are bounded by the vehicle's charging or discharging power, remaining battery capacity, available battery energy, and the duration of one tick.

C.3.2. Non-V2G charging

Drivers with access to a home charger can charge at home whenever their battery is not full. Away from home, they seek a charger only when the estimated state of charge after returning home would fall below 20%. Drivers without access to a home charger are more dependent on public and semi-public infrastructure and seek away-from-home charging when their state of charge is below 80%. Once a vehicle has obtained access to a charger, the assigned charging policy determines the requested charging power.

Normal charging

Normal charging applies the vehicle's maximum available charging power. The energy added during a tick is limited by the configured charging speed and the remaining battery capacity. Charging ends when the battery is full, when the vehicle departs, or when access to the charging point ends.

Net-aware charging

Net-aware charging applies a predetermined time-of-day power profile rather than responding to the current energy balance or net load. Between 00:00 and 06:00, the available charging power increases gradually from 5 kW to the maximum charging power of 10 kW. The maximum power remains available between 06:00 and 17:00, after which it is limited to 4 kW until 23:00. From 23:00 to midnight, the available power begins to increase again. The profile therefore reduces charging during the evening peak without requiring real-time information or forecasts.

Smart charging

Smart charging uses a rule-based calculation of the power needed to reach 80% state of charge before the next expected departure. For each eligible vehicle, the model determines the remaining energy requirement and the number of ticks until departure. Dividing these quantities gives the minimum charging rate required to reach the target on time, bounded by the maximum charging power and remaining battery capacity.

When additional charging capacity is available, part of this capacity may be used according to a grid-room factor derived from the current energy balance. The factor is bounded between zero and one

and is scaled using a fixed reference of 500 kW. Smart charging uses 40% of the available additional capacity, preventing it from consuming all available grid room.

A separate peak-avoidance rule stops charging when current net load exceeds 550 kW, but only when the target can still be reached after delaying charging by at least one tick. Consequently, the mobility requirement takes precedence when further delay would prevent the vehicle from reaching the 80% target before departure.

C.3.3. V2G control

The V2G controller considers participating vehicles that are located at their destination and connected to an available charging point. The term internal energy-balance signal is used here for the model variable previously described as a price/grid signal. This signal is derived from the simulated energy balance and does not represent an observed or forecast electricity-market price.

V2G control is executed sequentially. Emergency charging is applied first, after which the current system state is updated. Flexible charging is then allocated to the remaining eligible vehicles. Finally, discharging is evaluated using the net load remaining after both charging stages.

Emergency charging

Connected V2G vehicles whose battery energy is below their minimum battery requirement receive emergency charging. This minimum corresponds to the guaranteed range specified in the offered contract. Emergency charging is applied regardless of the energy-balance signal or current peak pressure.

The charging quantity is bounded by the energy required to restore the minimum battery level, the maximum charging power per tick, and the remaining battery capacity. The resulting demand is added to the system state before flexible charging and discharging are evaluated.

Flexible V2G charging

Flexible charging considers connected V2G vehicles that have not received emergency charging and whose batteries are not full. Whether charging is desirable is first determined using the internal energy-balance signal. The signal allows charging during renewable surplus and retains a moderate positive response close to balance so that participating vehicles do not remain permanently near their guaranteed minimum level. Within the near-balance range of approximately ± 10 kW, the signal is fixed at 0.40 to permit limited maintenance charging.

Three fixed parameters shape this signal. The soft-deficit setting allows limited maintenance charging during a mild deficit, with the signal decreasing approximately linearly from 0.40 near balance to zero at a residual energy balance of -350 kW. The strong-deficit setting determines how the retained negative branch develops under larger deficits, reaching its minimum at approximately -800 kW, while the surplus scale determines how strongly charging responds to renewable surplus, with the response increasing quadratically towards its maximum at $+400$ kW. In the final controller, only positive values permit flexible charging. Negative values therefore prevent flexible charging but do not trigger discharging, which is controlled separately by current net-load pressure. The resulting signal is illustrated in Figure C.1(a).

Flexible charging is additionally constrained by a soft peak guard. Let L_t denote current net load in kW. The final charging factor is:

$$f_{\text{charge}}(L_t) = \begin{cases} 1, & L_t \leq 450, \\ 1 - \left(\frac{L_t - 450}{700 - 450} \right)^2, & 450 < L_t < 700, \\ 0, & L_t \geq 700. \end{cases}$$

The factor therefore leaves flexible charging unrestricted below 450 kW, progressively reduces it between 450 and 700 kW, and stops it at 700 kW. The quadratic form retains most of the charging response near the start of the control range but reduces it more strongly as net load approaches the charging-stop threshold. This response is shown in Figure C.1(b).

When both the energy-balance signal and peak guard permit charging, the controller calculates a collective charging budget. This budget is limited by the system-level charging opportunity and the combined technical charging capacity of the eligible vehicles. It is distributed using weights that increase with the gap between the current battery level and the fixed V2G target state of charge of 80%.

V2G discharging

Discharging is evaluated after emergency and flexible charging. Eligible vehicles must not have charged in the same tick and must retain battery energy above the minimum level associated with their guaranteed range.

Discharging is triggered by current net-load pressure rather than by the internal energy-balance signal. This avoids using V2G energy for every small system deficit and preserves available battery energy for periods of greater peak pressure. The discharge-response factor is:

$$f_{\text{discharge}}(L_t) = \begin{cases} 0, & L_t \leq 550, \\ 0.85 \left(\frac{L_t - 550}{700 - 550} \right)^2, & 550 < L_t < 700, \\ 0.85, & L_t \geq 700. \end{cases}$$

The normalised net-load increase between 550 and 700 kW defines a linear peak-pressure signal between zero and one. The discharge response applies a quadratic transformation to this signal and is bounded at 0.85. Discharging therefore increases gradually after the start threshold and becomes stronger as net load approaches 700 kW, without using the full theoretical response. The two curves are shown in Figure C.1(c).

The collective discharge requirement is based only on the part of current net load above 550 kW. The controller therefore targets peak-relevant demand rather than attempting to reduce total net load to zero. Before technical vehicle constraints are applied, the responsive system-level discharge budget can be expressed as:

$$B_{\text{discharge}}(L_t) = \max(L_t - 550, 0) f_{\text{discharge}}(L_t).$$

This quantity is expressed as a power equivalent in Figure C.1(d). During model execution, it is divided by 12 to convert kW into kWh per five-minute tick. The resulting budget is subsequently limited by the combined technical discharging capacity and the usable battery energy available above the guaranteed minimum.

The budget is distributed according to usable battery energy. Vehicles with more energy above their minimum level receive a larger share, while vehicles close to the minimum contribute little or nothing. The discharged energy reduces current net load and is recorded separately by charging location.

Table C.4: Charging-control settings and their treatment in the experiments.

Procedure	Setting	Function	Treatment and final value
All charging	Maximum charging and discharging power	Technical power limit for each vehicle.	Fixed: 10 kW
Away-from-home charging	SOC threshold without home charger	Determines when a driver without home charging seeks another charger.	Fixed: 80%
Away-from-home charging	Return-home SOC threshold	Determines when a driver with home charging seeks an away-from-home charger.	Fixed: 20%

Continued on next page

Table C.4 – continued from previous page

Procedure	Setting	Function	Treatment and final value
Net-aware charging	Daily charging profile	Limits available charging power by time of day.	Fixed profile
Smart charging	Target state of charge	Battery target to be reached before the next departure.	Fixed: 80%
Smart charging	Grid-room scale	Scales the additional charging response to the current energy balance.	Fixed: 500 kW
Smart charging	Additional-capacity share	Fraction of spare charging capacity available for grid-responsive extra charging.	Fixed: 0.40
Smart charging	Peak-avoidance threshold	Delays charging above this net load when the target remains achievable.	Fixed: 550 kW
V2G charging	charge-reduce-start-kW	Net load at which flexible V2G charging begins to be reduced.	Explored; fixed at 450 kW
V2G charging	charge-stop-kW	Net load at which flexible V2G charging stops.	Explored; fixed at 700 kW
V2G discharging	discharge-start-kW	Net load at which flexible V2G discharging begins.	Explored; fixed at 550 kW
V2G discharging	full-discharge-kW	Net load at which the maximum discharge response is reached.	Explored; fixed at 700 kW
V2G charging signal	soft-deficit-kW	Permits maintenance charging during a mild system deficit.	Fixed: 350 kW
V2G charging signal	strong-deficit-kW	Scales the negative part of the internal energy-balance signal.	Fixed: 450 kW
V2G charging signal	surplus-scale-kW	Scales the charging response to renewable surplus.	Fixed: 400 kW
V2G flexible charging	Base minimum and maximum charging response	Bounds the fraction of the available collective charging budget, before scaling by the soft peak-guard factor.	Fixed: 0.35–1.00
V2G discharging	Minimum and maximum discharge response	Bounds the response applied to peak-relevant discharge demand.	Fixed: 0.00–0.85

Notes: “Explored” indicates that the setting was varied during the exploratory phase and subsequently fixed for final learning. The tested exploratory values are reported in Table B.1. The remaining settings were fixed model-design choices and were not dimensions of the final parameter search.

C.4. Detailed behavioural feedback and participation

This section provides the detailed preference, feedback, and contract-re-evaluation procedures underlying the summary in Section 4.6. It describes the latent preference classes, the timing and smoothing of psychological updates, and the way these mechanisms influence subsequent V2G participation.




Figure C.1: Response functions used by the V2G controller. Panel (a) shows the internal energy-balance signal used to determine whether and how strongly flexible charging is desirable; negative values prevent flexible charging but do not trigger discharging. Panel (b) shows the soft peak guard, which progressively restricts flexible charging between 450 and 700 kW.

Panel (c) distinguishes the linear peak-pressure signal from the bounded quadratic discharge response between 550 and 700 kW. Panel (d) shows the resulting responsive system-level discharge budget before vehicle-level technical constraints are applied. The budget is shown as a power equivalent and is converted to energy per five-minute tick during model execution.

C.4.1. Latent preference classes and contract evaluation

Driver preferences are based on the stated-choice experiment conducted by Bakhuis et al. (2025) among 1,018 Dutch respondents. The experiment included four V2G contract attributes: financial compensation, guaranteed driving range, minimum plug-in time, and additional battery degradation. A latent class choice model identified four groups with different sensitivities to these attributes.

Each EV-driver agent is assigned to a latent class according to the estimated population shares. The classes include two groups primarily concerned with guaranteed range, a group primarily concerned with plug-in flexibility, and a group primarily motivated by financial compensation. Table C.5 reports the class shares and the relative importance of the four contract attributes.

Within each class, individual preference coefficients are drawn around the estimated class means. This introduces heterogeneity both between and within latent classes. The resulting coefficients are used to calculate the utility of the offered contract. Participation occurs when this utility exceeds the driver-specific effective participation threshold.

Table C.5: Relative importance (RI%) of V2G contract attributes per latent class (Bakhuis et al., 2025)

Attribute	Class 1 (43%) Range-focused, finance-disregarding	Class 2 (29%) Range-focused, plug-in-time-disregarding	Class 3 (18%) Flexibility-focused users	Class 4 (10%) Finance-focused, range-disregarding
Financial compensation (FC)	9%	23%	11%	70%
Guaranteed range (GR)	66%	45%	25%	2%
Minimum plug-in time (MP)	12%	9%	53%	20%
Battery degradation (BD)	13%	23%	11%	8%

C.4.2. Update timing and inertia

Range anxiety, plug-in time concern, environmental self-identity, and, when enabled, peer influence are updated once per simulated day. The resulting values affect the driver-specific contract coefficients or participation threshold but do not immediately force a change in participation. Drivers re-evaluate the offered contract once per simulated week.

Changes in the psychological states are smoothed using an inertia mechanism. Let $x_{i,t}$ denote the current psychological state of driver i and let $\tilde{x}_{i,t+1}$ denote the value implied by the most recent experience. The updated state is:

$$x_{i,t+1} = 0.9x_{i,t} + 0.1\tilde{x}_{i,t+1}.$$

This prevents individual experiences from producing immediate large changes in contract preferences. Range anxiety and plug-in time concern are bounded between 0 and 1, while environmental self-identity is bounded between -1 and 1.

C.4.3. Range anxiety

Range anxiety represents concern about having insufficient battery energy for travel. It is updated after driving and after the corresponding battery consumption has been applied. An empty-driving event produces the strongest increase. Driving with a state of charge of 10% or lower produces a smaller increase, while completing travel with sufficient battery energy gradually reduces range anxiety.

In the preference update, range anxiety modifies the driver-specific sensitivity to guaranteed range. Higher range anxiety increases the value attached to contracts that preserve a larger minimum battery level. Negative mobility experiences therefore make the driver more sensitive to the mobility protection provided by the contract.

C.4.4. Plug-in time concern

Plug-in time concern represents the perceived inconvenience of remaining connected for the duration required by a V2G contract. The model accumulates plug-in time between 07:00 and 19:00 and evaluates the realised duration once per simulated day.

For V2G participants, realised plug-in time is compared with the minimum plug-in time specified in the offered contract. Plug-in concern increases when the requirement is not met and decreases when it is met or exceeded. For non-participating drivers, a baseline duration of eight hours is used as a reference for general plug-in experience.

Higher plug-in concern increases the disutility associated with the minimum plug-in-time attribute. Drivers who repeatedly experience the required connection duration as difficult to achieve therefore become less likely to accept contracts with long plug-in requirements.

C.4.5. Environmental self-identity

Environmental self-identity represents the extent to which a driver views their behaviour as environmentally or system supportive. The update is applied only when the vehicle has charged or discharged during the current tick and depends on how that realised action is classified by the model.

For non-V2G vehicles, charging is classified as grid-supportive when the energy balance before charging indicates sufficient renewable surplus. Charging close to balance or during a deficit is not classified as supportive. For V2G vehicles, emergency charging during a deficit is treated as non-supportive because it increases demand while restoring the contractual minimum battery level. Flexible V2G charging and V2G discharging are classified as supportive under the implemented congestion-oriented controller.

Higher environmental self-identity lowers the effective participation threshold, making V2G participation more likely. Repeated actions classified as non-supportive can reduce environmental self-identity and thereby weaken future willingness to participate.

C.4.6. Peer influence

The model retains an optional peer-influence mechanism. When enabled, the effect depends on both the number of peers participating in V2G and their apparent satisfaction. Peer satisfaction is approximated by the adoption margin, defined as the difference between a participating driver's contract utility and participation threshold.

A larger positive adoption margin represents a more favourable participation experience. More numerous and more satisfied participating peers lower the effective participation threshold, whereas less favourable peer experiences can increase it. The adjustment is bounded to prevent social influence from dominating the other participation mechanisms.

Peer influence was disabled in the final analysed experiments. The mechanism is documented here because it remains part of the model structure and was considered during model development.

C.4.7. Preference updating and contract re-evaluation

The updated psychological states affect contract evaluation through different pathways. Range anxiety changes the sensitivity to guaranteed range, plug-in time concern changes the sensitivity to minimum plug-in time, and environmental self-identity and peer influence affect the effective participation threshold.

Once per simulated week, each driver recalculates the utility of the offered V2G contract using the current preference coefficients. Drivers participate when contract utility exceeds their effective threshold; otherwise, they remain or become non-participants. The participation decision can therefore change during a simulation even though the offered contract attributes remain fixed within that run.

D

Model verification and validation

This appendix provides the supporting diagnostics and validation comparisons underlying the results summarised in Section 5.1. The verification and validation procedures themselves are described in Section 3.4. Results concerning loss-function selection, stochastic variation, and the number of repetitions are reported separately as part of the exploratory analysis.

D.1. Verification diagnostics

Table D.1 reports the principal numerical and logical checks applied to the final model version. It presents both the criterion evaluated and the outcome of the corresponding diagnostic runs.

Table D.1: Numerical and logical diagnostics used for model verification.

Diagnostic	Verification criterion	Result in final diagnostic runs
Energy accounting	Reported net load can be reconstructed from residential and service-sector demand, renewable generation, EV charging, and V2G discharging.	The reconstructed and reported balances were equal; both the balance error and split-net error were zero.
Battery boundaries	No vehicle exceeds its battery capacity or falls below its applicable minimum battery level.	No violations were identified in the final diagnostic runs.
Applied charging power	Applied charging and discharging setpoints remain within the technical limits and available battery headroom.	Requested setpoints were limited correctly when the full requested power could not be applied.
Simultaneous actions	Vehicles cannot charge and discharge simultaneously and cannot drive while charging or discharging.	No simultaneous charging and discharging or movement-and-charging cases were detected.
Charging location	Charging or discharging can occur only when the vehicle is located at a patch containing a suitable charging point.	All inspected charging and discharging actions occurred at valid charging locations.
Guaranteed range	V2G discharging cannot reduce the battery below the minimum level implied by the guaranteed range.	No discharge below the applicable battery minimum was detected after correction of the discharge procedure. Higher guaranteed-range settings increased the protected battery reserve and reduced available discharge headroom.
Driving-distance classification	Each distance increment is counted exactly once after battery use: as empty-battery driving when the battery reaches zero, low-battery driving when state of charge is at most 10%, or regular electric driving otherwise. V2G and non-V2G subgroup counts must sum to the corresponding totals.	Battery energy was deducted before classification. Empty, low, and regular electric kilometres were mutually exclusive, their sum equalled the total distance driven, and subgroup totals were consistent with the aggregate counters.
Charger allocation	Reservations and occupancy states are created and released consistently with vehicle eligibility and actual charger use.	Large-scale unused occupancy was removed after adjustment of the allocation rule, and no stale reservations were detected in the final checks.
Smart-charging feasibility	Charging is postponed only when the target battery level can still be reached before the next expected departure.	The required energy, remaining connection time, and achievable charging volume produced the expected charging decisions.
V2G control and model integration	The sequence of charging and V2G procedures, their effect on the energy balance, and their interaction with vehicle-level battery and power constraints follow the implemented control logic.	Tick-level traces showed that emergency and flexible charging updated the system state before discharging was evaluated. V2G discharging occurred during elevated net load when vehicles had sufficient usable battery energy, and vehicles did not charge and discharge within the same tick.
Participation and behavioural logic	Contract utility, latent-class assignment, participation thresholds, and psychological updates follow the implemented decision rules and remain within their specified bounds.	Selected agent traces showed consistent utility-based participation decisions and correctly bounded behavioural updates.
Experimental pipeline	Parameter values, random seeds, reporter outputs, repetition identifiers, and calculated losses remain complete and traceable for every stored simulation.	Final observations retained the required experimental metadata. Incomplete, unreadable, or invalid outputs were identified before analysis.

D.2. Supporting validation comparisons

The following comparisons provide additional detail on the empirical and qualitative validation results summarised in the main text. The comparisons concern the timing and shape of the simulated profiles rather than exact numerical correspondence, because the modelled population, infrastructure, control strategies, and aggregation methods differ from the available reference data.

D.2.1. Location-specific charging profiles

The temporal charging patterns were compared qualitatively with the location-specific profiles reported by ElaadNL (Refa et al., 2023). Across the different charging locations, the simulated profiles showed a comparable development over the day.

Residential charging increased strongly after vehicles returned home and reached its highest level during the evening. The simulated evening peak was somewhat higher than in the ElaadNL profile, which is likely related to the relatively high and homogeneous charging power used in the model. Because more energy was charged shortly after arrival, the simulated profile also declined somewhat faster during the night.

The simulated net-aware profile corresponded particularly well with the temporal development of the net-aware public-charging profile reported by ElaadNL. Both showed reduced charging during the constrained evening period and a subsequent recovery when more charging capacity became available. The morning peak was somewhat smaller in the model, which may be related to the aggregation of weekdays and weekend days, whereas the ElaadNL reference profile represents weekdays only.

The profiles therefore do not provide an exact empirical fit, but their timing, shape, and response to charging constraints were sufficiently similar to support the plausibility of the implemented charging behaviour.

D.2.2. Mobility outcomes

Mobility patterns were examined as a plausibility check. In an eight-week diagnostic run, the simulated vehicles travelled an average of 14.58 km per vehicle per day. This is of a similar order of magnitude to the 15.14 km per car driver per day reported in table G.4 (Centraal Bureau voor de Statistiek, 2016). The measures are not fully equivalent because the simulation result is expressed per vehicle, whereas the reference is expressed per driver. The comparison was therefore used only as a broad plausibility check.

Across the validation runs, the share of kilometres driven while vehicles had no remaining battery energy remained below 0.239% in 95% of the runs. A small number of outlying runs showed higher values, with a maximum of 1.915%. These outcomes indicate that the charging and discharging controls generally preserved vehicle availability, although unfavourable configurations could still produce occasional mobility shortages.

D.2.3. Smart-charging and V2G profiles

Smart charging and V2G are not yet deployed at a scale that provides a single representative empirical charging profile. Their temporal behaviour also depends strongly on the intended service and associated control strategy. Whereas many emerging applications are primarily driven by electricity prices, the smart-charging and V2G controls in this study were designed to reduce local congestion pressure.

The resulting profiles were therefore compared qualitatively with the charging scenarios reported by ElaadNL (ElaadNL, Team Marktonontwikkeling, 2025). For V2G, the magnitude of charging and, particularly, discharging peaks varied with the selected control thresholds and available participation. The general pattern nevertheless remained consistent: charging was shifted away from constrained periods where possible, while discharging was concentrated during elevated evening net load.

The smart-charging profile was flatter than the uncontrolled and V2G profiles. Charging was reduced during the evening peak but did not always stop completely because vehicles still required sufficient time to reach their battery target before departure. The morning charging peak was therefore smaller and more dispersed. This was partly caused by the use of expected rather than perfectly known departure times, which allowed many vehicles to distribute their required charging more gradually over the

available connection period.

These comparisons were not interpreted as evidence that the model reproduces one expected future charging profile exactly. Instead, they indicated that the implemented control strategies generated temporal responses consistent with their intended congestion-oriented functions.

D.2.4. Expert review

The simulated charging profiles were additionally discussed with Jolijn van Dijk of EVConsult. The review focused on whether the aggregated temporal patterns generated by normal, net-aware, smart, and V2G charging were recognisable and plausible in practice.

The discussion considered charging power, connection duration, battery requirements, the timing of controlled charging, and the additional contribution of V2G discharging relative to unidirectional smart charging. The review supported the use of the simulated profiles for comparing charging strategies, while emphasising that real-world outcomes depend on local network conditions, charger settings, control objectives, and user participation.

E

Exploratory results

This appendix provides the supporting results of the exploratory analysis summarised in Section 5.2. It documents the experiment series, the assessment of alternative congestion-loss definitions, and the exploratory learning and stability analyses used to define the final experiments.

E.1. Runs

Table E.1 summarises the experiment series included in the exploratory analysis, together with their number of runs, repetitions, parameter combinations, and analytical purpose. The series differ in scope and repetition count because they were designed for different stages of exploration, sampler testing, and stability assessment.

The Bayesian series also differed in the amount of training data available at their start. The BVM series was initialised using results from the three preceding broad exploratory series, vOy, U1s, and b3Q. For the subsequent trained Bayesian series, the training data were accumulated: each later series used all the available results from the preceding broad and corrected series. The final stability series used the cumulatively expanded training set, including vOy, U1s, b3Q, BVM, hjM, DJ2, and 0ZD.

Table E.1: Overview of experiment series included in the exploratory analysis

Experiment series	Runs	Reps.	Combinations	Purpose and configuration
jT0_run	40	5	8	Initial exploration of V2G control inputs using a reduced parameter space.
bOg_srun1	200	50	4	Initial repetition-stability test using V2G control inputs and a restricted parameter space.
SC5_srun2	200	50	4	Parallel repetition-stability test using V2G control inputs and a restricted parameter space.
kDF_run1	140	5	28	Broad exploration of the parameter space, including V2G control inputs.
dJV_run2	140	5	28	Parallel broad exploration using the same general parameter space.
v55_run3	240	5	48	Additional broad exploration of the parameter space.
SXu_run1	300	3	100	Broad parameter-space exploration with three repetitions per combination.
DcZ_run2	300	3	100	Parallel broad parameter-space exploration.

Continued on next page

Table continued from previous page

Experiment series	Runs	Reps.	Combinations	Purpose and configuration
yVF_run3	300	3	100	Parallel broad parameter-space exploration.
vOy_run1	300	3	100	Further exploratory Bayesian series using the broad parameter space.
U1s_run2	300	3	100	Parallel exploratory Bayesian series using the broad parameter space.
BVM_run1	250	5	50	First series after correction of the Bayesian sampling procedure, preloaded with results from preceding exploratory series.
hJM_run2	300	3	100	Corrected Bayesian sampler using an exploitative sampling strategy.
DJ2_run1	250	5	50	Corrected Bayesian sampler using the dynamic sampling strategy.
OZD_run2	250	5	50	Parallel corrected Bayesian series using the dynamic sampling strategy.
vOZ_srun1	250	50	5	Final stability series using the corrected, exploitative sampler preloaded with the preceding corrected experiment results.
dKy_srun2	250	50	5	Parallel final stability series using the same corrected sampler configuration.
oS8_srun3	250	50	5	Additional final stability series using the corrected sampler configuration.
GXU_srun4	250	50	5	Additional final stability series using the corrected sampler configuration.
Total	4,510		890	

E.2. Loss function

This appendix reports the supporting diagnostics for the loss-function assessment summarised in Section 5.2.2. These include the distribution of alternative losses, their rank correlations, and their stochastic variability.

E.2.1. Overview tested losses

Table E.2 summarises the distribution of the squared-exceedance loss across alternative thresholds. Higher thresholds produce more zero-loss runs and therefore less distinction between parameter combinations; the 600 kW threshold retains substantially more non-zero variation than the 650 and 700 kW alternatives.

E.2.2. Rank-correlation loss

Figure E.1 shows Spearman rank correlations for a representative selection of the tested loss definitions. The 500 and 600 kW formulations generally produce similar rankings, whereas the 700 kW formulations correlate strongly with one another but less strongly with the lower-threshold losses.

E.2.3. Stability loss

Figure E.2 compares the median relative MAD across the tested thresholds and penalty forms. Count-based, linear, squared, and cubic losses show comparatively limited relative variability, whereas the tenth- and fiftieth-power penalties show substantially higher values. The 600 kW squared loss remains within the moderate range of the tested formulations.

Table E.2: Descriptive comparison of the squared congestion loss across alternative thresholds.

Threshold (kW)	Zero loss (%)	Median	IQR	95th percentile	Maximum
400	0.0	1.285×10^3	6.024×10^2	3.466×10^3	1.435×10^4
450	0.0	4.150×10^2	2.391×10^2	1.445×10^3	8.593×10^3
500	0.0	9.910×10^1	7.303×10^1	5.312×10^2	5.186×10^3
550	0.0	1.324×10^1	1.394×10^1	1.532×10^2	3.131×10^3
600	0.0	6.237×10^{-1}	1.074	4.434×10^1	1.876×10^3
650	83.9	0	0	1.872×10^1	1.106×10^3
700	84.2	0	0	7.835	6.327×10^2

Note. Statistics are calculated across individual simulation runs using the squared exceedance loss. Zero loss indicates that the simulated net load did not exceed the corresponding threshold during the run. The 600 kW threshold is highlighted because it was selected




Figure E.1: Spearman rank correlations between a representative selection of congestion-loss definitions. The selected definitions combine exceedance counts and linear, squared, and cubic excess penalties at thresholds of 500, 600, and 700 kW and were chosen to illustrate the broader correlation pattern observed across the tested loss formulations. Higher correlations indicate that parameter combinations are ranked similarly.




Figure E.2: Relative stochastic variability of count- and power-based congestion-loss definitions across repeated simulation runs. Relative variability is expressed as the median absolute deviation divided by the median loss for each parameter combination. Values are shown for thresholds between 450 and 800 kW and for exceedance count and linear, squared, cubic, tenth-power, and fiftieth-power penalties.

E.3. Exploratory learning

Table E.3 summarises the distribution of individual-run loss within the experiment series used to examine the development of the exploratory search. The later corrected and stability series generally show lower median losses and narrower upper tails than the earlier broad-search series. Because the series differed in parameter space, purpose, and number of repetitions, the statistics provide a descriptive overview rather than a like-for-like comparison of sampler performance.

E.3.1. Stochastic variability of the exceedance-count loss

Figure E.3 compares the median exceedance-count loss with its relative variability across repetitions. No clear monotonic relationship is visible, indicating that the typical number of exceedance ticks provides limited information about the relative stochastic variability of a parameter combination.

Table E.4 reports the exploratory comparison of the V2G control thresholds and identifies the settings retained for the final experiments.

Table E.3: Summary statistics of individual-run loss by exploratory experiment series.

Experiment series	Runs	Min	Q25	Median	Q75	Max	MAD	Mean	SD
SXrun_1	250	0.622	2.583	14.953	97.927	196.066	13.634	48.547	57.169
Dcrun_2	300	0.393	3.747	53.350	127.527	1352.867	51.409	93.069	161.279
yFrun_3	276	0.374	4.903	11.346	51.413	3133.973	10.008	98.346	334.434
v0_run1	300	0.352	4.718	16.831	111.121	3035.239	15.384	199.106	519.508
U1_run2	300	0.237	4.367	27.799	128.874	1952.747	26.020	148.271	329.105
BVM_run	250	0.250	0.531	1.077	13.383	54.052	0.723	8.936	14.032
hJM_run	288	0.318	1.541	8.156	18.124	1875.730	6.851	48.943	181.245
DJ_run1	250	0.261	0.494	0.623	1.424	139.747	0.221	9.951	25.038
oZ_run2	250	0.248	0.447	0.741	1.416	193.114	0.385	8.934	29.587
vO_Z_run1	250	0.256	0.502	1.282	1.467	4.083	0.281	1.062	0.580
dKy_run2	250	0.219	0.319	0.396	0.499	0.851	0.084	0.419	0.124
oS8_run3	250	0.220	0.348	0.419	0.492	0.774	0.073	0.427	0.103
GXU_run4	250	0.307	0.495	0.642	0.773	9.296	0.146	1.015	1.214

Note: Statistics are calculated across individual model runs within each experiment series and are not aggregated by unique parameter combination. Q25 and Q75 denote the first and third quartiles; MAD denotes the median absolute deviation.




Figure E.3: Relationship between the median exceedance-count loss and relative stochastic variability at the 500 kW threshold. Each point represents one parameter combination. The horizontal axis shows the median number of exceedance ticks across repetitions, while the vertical axis shows the corresponding MAD divided by the median count; both axes are logarithmic. No clear monotonic relationship is visible between the typical exceedance count and its relative variability.

Table E.4: Exploratory comparison of the V2G control thresholds. Bold values indicate the settings retained for the final experiments.

Control setting	Thr.	Median loss	Mean loss	Min. loss	Max. loss	Median rep. SD	Driven-empty V2G
Charge reduction	400	72.69	933.30	0.00	2741.37	11.46	467.36
	450	30.12	215.32	14.42	601.43	7.54	198.51
	500	15.88	48.77	0.13	262.93	2.61	1138.16
Charge stop	600	2.12	2.45	0.13	5.10	0.27	1725.12
	650	432.18	432.18	262.93	601.43	25.38	177.12
	700	47.88	486.73	15.72	1835.42	7.56	515.12
	750	17.04	412.01	0.00	2741.37	4.05	635.36
Discharge start	500	262.93	734.27	0.00	2741.37	21.46	647.84
	550	17.04	17.10	0.13	30.12	4.05	583.01
	600	15.23	324.30	2.12	1835.42	2.48	980.56
Full discharge	600	47.88	287.97	0.00	1835.42	7.56	668.76
	650	2.12	5.99	0.13	15.72	0.27	2039.20
	700	1373.24	1373.24	5.10	2741.37	3020.72	29.12
	750	17.04	211.50	16.04	601.43	5.39	170.56

F

Final results

This appendix provides supplementary results supporting the sampler-learning, parameter-importance, seasonal-comparison, and PRIM analyses presented in Section 5.3. The results are reported separately to provide the detailed figures and diagnostics underlying the main findings.

Table F.1 provides an overview of the proposal-level loss statistics for the final Bayesian experiment series, reported separately for the two learning chains.

Table F.1: Overview of proposal-level loss statistics for the final Bayesian experiment series

Series	Min	Max	Median	MAD	Unique
V2G_1					
lep_run1A	0.282	303.445	2.840	2.451	20
Q7n_run1B	0.399	564.344	2.109	1.678	20
CzW_run2A	0.232	116.537	0.839	0.478	22
r3d_run2B	0.260	18.392	1.000	0.580	22
AnP_runA	0.247	137.598	0.550	0.203	22
Emk_run3B	0.239	10.880	0.589	0.253	22
PWQ_run4A	0.219	17.428	1.539	1.178	22
irw_run4B	0.260	6.800	0.527	0.192	22
bjJ_run5A	0.256	32.815	0.650	0.285	22
vpo_run5B	0.254	5.652	0.811	0.456	22
Cel_run6A	0.241	10.517	0.547	0.196	22
rMe_run6B	0.234	32.480	0.645	0.341	22
Lil_run7A	0.244	5.178	0.590	0.159	22
dQY_run7B	0.234	5.534	0.708	0.334	22
V2G_2					
QAQ_run1A	0.423	84.361	1.825	1.060	20
XEt_run1B	0.311	60.493	3.744	3.109	20
QgZ_run2A	0.286	8.579	0.656	0.243	24
hXr_run2B	0.263	3.463	0.851	0.386	24
7Mn_run3A	0.270	419.407	0.648	0.275	22
QIL_run3B	0.290	13.052	0.644	0.286	22
Qpe_run4A	0.243	26.928	0.470	0.114	22
FxH_run4B	0.239	55.718	0.517	0.132	22
0Qr_run5A	0.237	1.725	0.540	0.146	22
soQ_run5B	0.270	30.612	0.610	0.219	22
Clc_run6A	0.251	1.431	0.532	0.119	22
W4r_run6B	0.223	2.211	0.633	0.265	22
P7J_run7A	0.262	2.292	0.555	0.178	22
RQZ_run7B	0.241	1.237	0.526	0.111	22




Figure F.1: Proposal-level median loss across internal Bayesian batches for selected learning rounds (1, 3, and 4) of V2G 1 (top row) and V2G 2 (bottom row). Points represent the median loss of individual parameter proposals across repetitions, while dashed lines show the median proposal loss within each internal batch. Each colour denotes one of the two parallel runs conducted in a learning round. Loss is displayed on a logarithmic scale.

F.1. Sampler learning

This section provides the supporting results for the sampler-learning analysis presented in Chapter 5. It reports changes between successive learning rounds and within the internal Bayesian batches of individual rounds. Results are shown separately for the two learning chains.

F.1.1. Between-round learning

Table F.2 reports the development of proposal-level loss between successive learning rounds. The quartile changes describe shifts in the distribution of evaluated proposals relative to the preceding round.

F.1.2. Within-round learning

Table F.3 compares proposal loss between the first and final internal Bayesian batches within each learning round. Figure F.1 shows the corresponding proposal-level trajectories for selected rounds in both learning chains.

F.2. Feature scores rounds

Figures F.2 and F.3 report the development of parameter feature scores across learning rounds and their robustness across alternative model specifications and scoring methods. Figure F.5 compares the predictive performance of Extra Trees models fitted to raw and log-transformed congestion loss and loss variability.

Table F.2: Median proposal loss and round-to-round changes in the quartiles of proposal-level loss. Positive change values indicate lower loss relative to the preceding learning round; negative values indicate deterioration.

Learning chain	Round	Median loss	Q25 change (%)	Median change (%)	Q75 change (%)
V2G_1	1	2.835	—	—	—
V2G_1	2	0.978	49.5	65.5	93.9
V2G_1	3	0.593	13.8	39.4	1.5
V2G_1	4	0.728	-9.3	-22.7	9.9
V2G_1	5	0.708	8.2	2.7	27.9
V2G_1	6	0.593	-5.0	16.2	-8.0
V2G_1	7	0.643	-8.3	-8.4	17.3
V2G_2	1	2.417	—	—	—
V2G_2	2	0.719	47.2	70.3	85.5
V2G_2	3	0.612	14.9	14.8	-9.0
V2G_2	4	0.510	7.7	16.6	49.9
V2G_2	5	0.561	-6.6	-9.9	-32.2
V2G_2	6	0.573	-2.5	-2.1	1.4
V2G_2	7	0.544	-1.3	5.1	12.1

Note. Each parameter configuration was first summarised across its repetitions. Median loss is the median proposal-level loss in the corresponding learning round. Quartile changes were calculated relative to the corresponding quartile in the preceding round as $100 \times (q_{r-1} - q_r) / q_{r-1}$. Positive values therefore indicate a reduction in loss, whereas negative values indicate an increase.

Table F.3: Percentage change in median proposal loss from the first two to the last two internal Bayesian batches within each learning round. Positive values indicate lower loss in the final batches; negative values indicate higher loss.

Learning round	V2G_1 change (%)	V2G_2 change (%)
1	67.0	79.9
2	8.8	50.7 ^a
3	79.9	-4.1
4	-61.2	31.0
5	20.4	3.7
6	-49.6	-14.6
7	-9.6	-28.6

Note. Change was calculated as $100 \times (\text{early median loss} - \text{late median loss}) / \text{early median loss}$, where early and late refer to the first two and last two internal Bayesian batches, respectively. Positive values indicate a reduction in median proposal loss; negative values indicate deterioration. ^aOne V2G_2 experiment series in round 2 reached the computational time limit after six complete internal batches.



Figure F.2: Median normalised Extra Trees feature scores for log-transformed congestion loss across successive learning rounds, with both V2G model chains combined. The corresponding out-of-bag R^2 values averaged approximately 0.72.




Figure F.3: Comparison of normalised parameter feature scores across alternative models, target definitions, and scoring methods. The figure is used as a robustness check on the ranking obtained with Extra Trees and log-transformed loss.




Figure F.4: Parameter consistency among the five lowest-loss configurations in each learning round. Each cell shows the share of the five configurations taking the most common value of that parameter. High consistency indicates convergence towards a common parameter value among the best-performing proposals, but does not imply causal importance or convergence to a unique global optimum.

F.3. Parameter repetition

Figure F.4 reports the consistency of parameter values among the five lowest-loss configurations in each learning round.

F.4. Seasonal performance comparison

Figures F.6 and F.7 compare model performance between winter and summer. They report the persistence of relative performance between seasons and the relationship between seasonal V2G participation and congestion-related loss.

F.5. PRIM

Figure F.8 reports the bootstrap recovery of the parameter restrictions defining the selected PRIM box. Figures F.9 and F.10 report the cross-chain validation results and the coverage–density trade-offs for the selected PRIM regions.




Figure F.5: Comparison of Extra Trees out-of-bag R^2 for models fitted to raw and log-transformed congestion loss and loss variability for the two V2G experiment groups.




Figure F.6: Seasonal performance transitions between winter and summer. Configurations are grouped into the best 20%, middle 60%, and worst 20% based on congestion-related loss in each season. Percentages indicate the share of each winter performance group that falls into the corresponding summer performance group.




Figure F.7: Relationship between mean seasonal V2G participation and congestion-related loss for the two learning chains. Each point represents an evaluated parameter configuration. Spearman rank correlations are reported in the legends. Congestion loss is shown on a logarithmic scale.




Figure F.8: Bootstrap recovery of the parameter restrictions in the selected PRIM box. Values show the percentage of bootstrap samples in which the corresponding restriction was recovered. The dashed line indicates an 80% recovery threshold.




Figure F.9: Cross-chain validation of the best-20% PRIM regions. Boxes derived from one learning chain were evaluated on the other chain using coverage and density. Filled markers indicate performance on the training chain and open markers performance on the independent test chain; lift values are shown separately.




Figure F.10: Coverage–density trade-offs along the PRIM peeling trajectories for the best-performing 10% and 20% targets. Each point represents a candidate PRIM box, with colour indicating the number of restricted dimensions.