Exploring Policy Interventions to Promote Equitable Electric Vehicle Adoption Using Agent-Based Modelling

EPA Master Thesis

T.L. Dert

Delft University of Technology

TU Delft

MSc thesis in EPA

Exploring Policy Interventions to Promote Equitable Electric Vehicle Adoption Using Agent-Based Modelling

Tamara Lisanne Dert
Student number: 5179440

Date: August 12, 2025
To be defended: August 29th, 15:00 at TPM Hall G

A thesis submitted to the Delft University of Technology in partial fulfilment of the requirements for the degree of Master of Science in Engineering & Policy Analyses

Tamara Lisanne Dert: *Exploring Policy Interventions to Promote Equitable Electric Vehicle Adoption Using Agent-Based Modelling* (2025)
Cover: AI Generated using Freepik Mystic 2.5 Flexible

The work in this thesis was carried out at:

Dep. Sustainable Urban Mobility & Safety
TNO, The Hague

In collaboration and submitted to:

Fac. Technology, Policy and Management
Delft University of Technology, Delft

First Supervisor TU Delft:	Prof.dr. Natalie van der Wall
Second Supervisor TU Delft:	Prof.dr. Tina Comes
TNO supervisor:	Dr.ir. Yashar Araghi

Executive Summary

This master's thesis investigates how policy interventions affect the speed and equity of electric vehicle (EV) adoption in the Netherlands, using an agent-based model (ABM) grounded in behavioural decision-making. The model simulates household-level adoption decisions between 2022 and 2035, drawing on the CODEC framework to reflect attention, enabling conditions, and intention formation. Socio-demographic diversity is represented through spatially clustered household archetypes based on income, education, and infrastructure access. The model is calibrated to reflect Dutch conditions and includes an explicit application to The Hague.

The study explores how different combinations of government interventions, including targeted subsidies, infrastructure investment, awareness campaigns, and zero-emission zones, can shape adoption trajectories across neighbourhoods. A key focus is the trade-off between accelerating overall EV uptake and ensuring an equitable transition across socio-economic contexts. In addition to literature-based policy scenarios, an exploratory modelling approach was used to generate and test a wide range of policy timing combinations under uncertainty.

The findings show that while comprehensive strategies (e.g. combining subsidies, marketing, and zero-emission zone regulation) perform best overall, they deliver only modest gains over simpler, well-timed interventions. Improvements of around 5 percentage points in EV share and moderate reductions in inequality are possible, but come with distinct implementation demands. Simpler strategies, such as infrastructure and marketing alone or early subsidies with infrastructure, often achieve comparable outcomes with less complexity.

Overall, the results highlight the importance of behavioural diversity, timing, and adaptability in policy design. A just and accelerated EV transition is feasible, but not automatic, and requires deliberate, strategically layered interventions. Achieving this requires planning further ahead and having adaptive responses ready for an uncertain future. This thesis contributes to the literature on sustainable mobility transitions by integrating behavioural realism, spatial equity, and exploratory policy design into a unified simulation framework.

Keywords: Agent-Based Modelling, Electric Vehicles, Equity in Mobility, Sustainable Mobility, Adoption Behaviour, Socioeconomic Barriers

Contents

1. Introduction **1**

2. Literature and Context **3**
- 2.1. Literature Review 3
  - 2.1.1. Determinants of EV Adoption 3
  - 2.1.2. Consumer Decision Process 6
  - 2.1.3. Equity in EV Policy Measures 9
- 2.2. Context 10
  - 2.2.1. EV Adoption in Europe and Beyond 10
  - 2.2.2. Policy Approaches and Effectiveness 10
  - 2.2.3. The Hague as a Case Study 11
- 2.3. Conclusion 13

3. Methods **14**
- 3.1. Overall Research Design 14
- 3.2. Methodological Approach 15
  - 3.2.1. ABM Modelling Approach 15
  - 3.2.2. Archetypical Approach 15
  - 3.2.3. Data Sources 16
  - 3.2.4. Research Scope 16

4. Model Conceptualization and Formalization **17**
- 4.1. Conceptual Framework 17
  - 4.1.1. Model Purpose 17
  - 4.1.2. Time and Space 17
  - 4.1.3. Overview 17
- 4.2. Formal Model 20
  - 4.2.1. System Initialization and Dynamics 21
  - 4.2.2. Agents: Household Initialization and behaviour 23
  - 4.2.3. Agent: Government Policy Mechanisms 27

5. Model Validation and Exploration **31**
- 5.1. Verification and Validation 31
  - 5.1.1. Agent-Level 31
  - 5.1.2. System-Level 32
  - 5.1.3. Policy-Level 33
- 5.2. Exploration 35
  - 5.2.1. Scenario and Uncertainty Exploration 35
  - 5.2.2. Policy Exploration 37
  - 5.2.3. Selected Policy Configurations for Further Exploration 38
  - 5.2.4. Experimental Design 40

6. Results: Robust Policy Evaluation under Deep Uncertainty **43**
6.1. Scenario-Dependent Policy Performance 43
6.1.1. Distribution of Policy Outcomes 44
6.1.2. Robustness Scoring 46
6.1.3. Scenario Discovery 47
6.2. Trade-off Analysis Between Objectives 50
6.2.1. Correlation of Outcomes 50
6.2.2. Parallel Coordinates Plot 51
6.2.3. Pareto Frontier and Objective Filtering 52
6.2.4. Statistical Significance of Policy Outcomes 53
6.3. Implementation in The Hague: From Strategy to Spatial Reality 55
6.3.1. EV Adoption Trajectories and District-Level Outcomes 55
6.3.2. Comparing Selected Policies in The Hague Simulation 56

7. Discussion **59**
7.1. The Main Findings 59
7.2. Discussion of Policy Implementation 60
7.3. Strengths and Limitations 61
7.4. Future Research Directions 62
7.5. Theoretical and Practical Implications 62

8. Conclusion **64**

A. Policy Significance Test **65**

B. PCA Policy Combination Clustering Results **66**

List of Figures

2.1. Decision diagram of the modeling method from Novizayanti et al. [2021]. . . . 8
2.2. Share of new cars sold that are electric, 2010–2024 [Ritchie, 2024]. 11
2.3. Income distribution of Dutch EV and non-EV drivers. Higher-income households are overrepresented among EV owners. Based on data of European Commission [2024] 12
2.4. (a) Spread of EV adoption levels in The Hague, 2024; (b) Spread of standardised income across The Hague, 2021 (both based on data from [Municipality The Hague, 2024]). 12

4.1. Overview of conceptual ABM setup, adapted from Huang et al. [2021]. 18
4.2. Determinants of agent behaviour in EV decision-making. 19
4.3. Pseudocode overview of the simulation model's operational flow. 20
4.4. (a) Elbow plot for cluster selection. (b) Radar plot of archetype features. (c) Spatial distribution across The Hague. 25

5.1. Playmodel Shot: Verification of Spatial Grid and Emergent Behaviour 32
5.2. (a) Simulated vs. Historical EV Adoption (2022–2025). (b) Simulated Baseline vs. Archetype Adoption Patterns (2022–2035) 33
5.3. Parameter-Wise and Combined Perturbation Effects (±10%) on EV Adoption Outcomes 34
5.4. Policy implementation 34
5.5. (a) SOBOL global sensitivity analyses convergence. (b) SOBOL global sensitivity analyses parameter outcomes. 36
5.6. EV Adoption under Policy Scenarios across Archetypes. Grey lines: individual policies; dashed line: base scenario. 37
5.7. Left: Adoption vs. inter-archetype variance. Right: Clustering of top-performing policy sequences. 38
5.8. Policy timing frequency among top- and bottom-performing clusters. 39

6.1. Distribution of key outcome metrics (1,000 scenarios per policy) across six policy configurations. 45
6.2. Robustness scores. left: share of scenarios meeting 0–4 outcome thresholds. right: percentage of scenarios per policy meeting each individual threshold. . 47
6.3. PRIM trade-off plots per policy, showing density versus coverage when identifying high-performing subsets of the uncertainty space. Color indicates the number of restricted dimensions. 48
6.4. Restricted ranges per uncertainty across all successful PRIM boxes. Median (bars) and full (black lines). 49
6.5. Correlation matrix showing relationships among key outcome metrics. 51
6.6. Parallel coordinates plot showing trade-offs patterns by averages across scenarios 52

6.7. Scatterplot matrix showing pairwise trade-offs between four normalized objectives for each policy combination, with colors indicating policy type and shapes denoting Pareto-optimality in the averaged four-objective space. 53
6.8. EV adoption over time by district (left) and predicted EV fleet percentage by 2038 (right). The black dashed line represents the weighted average across The Hague. 56
6.9. (a) Radar plot of archetype features. (b) Spatial distribution across The Hague. 56
6.10. EV adoption over time by district under *Infra+Marketing* (PC2) (upper) and under *Subsidy+Awareness+ZEZ* (PC4) (lower). 58

B.1. PCA clustering result 66

1. Introduction

Climate change is accelerating, leading to more extreme weather, food insecurity, health risks, and loss of biodiversity [European Commission, 2025]. Without timely intervention, these effects are projected to exacerbate global poverty and displacement [United Nations, 2025]. In response, the UN's Paris Agreement identifies the electrification of transportation as one of the key mitigation strategies, with electric vehicles (EVs) playing a central role in reducing carbon emissions [IEA, 2024]. The Netherlands has echoed this ambition in its 2019 Climate Agreement, which requires that all new cars sold after 2030 be emission-free [Ministerie van Algemene Zaken, 2019]. Despite these ambitions and the corresponding 58 million euros of allocated subsidies, projections suggest that conventional gasoline-powered cars will still dominate the Dutch market by 2030, with a share exceeding 70% [Paradies et al., 2023]. This gap between policy goals and market trends underscores the urgency of addressing barriers to EV adoption more effectively. In addition, considerable controversy has emerged regarding their equity [More, 2025]. Equity is, the fairness and inclusivity of the transition to electric mobility [Hopkins et al., 2023]. Although financial incentives and charging infrastructure have been heavily promoted, adoption has remained concentrated among wealthier households, leaving lower-income and marginalized communities behind. Current adoption patterns reflect these disparities and highlight the risk they pose to both the speed and fairness of the energy transition[European Commission, 2024].

Existing literature on EV adoption has highlighted the importance of factors such as economic incentives, infrastructure availability, and social influence [Hopkins et al., 2023; Sikder et al., 2023; Bhat and Verma, 2022]. Economic incentives lower the financial barrier to entry, infrastructure reduces range anxiety and increases convenience, and social influence shapes perceptions of EVs as desirable or normative within peer networks. However, most studies focus on one dimension at a time—often limited to new vehicle sales—and do not fully account for how these factors interact over time or across diverse population groups. This research builds on such prior work, particularly the projections by Paradies et al. [2023], but extends the scope by including second-hand markets and shared mobility options. These additions broaden the lens to consider a greater share of the fleet, and deepen the investigation into policy timing, socio-demographic heterogeneity, and evolving behavioural dynamics. This specific broadening is chosen because, current policy design often unintentionally reinforces inequality. Financial incentives tend to benefit wealthier households who already have access to capital, while lower-income and marginalized communities face structural barriers such as limited availability of public charging infrastructure, high upfront costs, and lower awareness [Lee et al., 2021; Song and Potoglou, 2020]. Equity in this context refers to whether different demographic groups can equally access the benefits of the EV transition, including the ability to afford and charge a vehicle. Addressing these disparities is essential not only from a justice point of view but also for effectiveness, as wide-scale adoption is crucial to reaching climate goals. Understanding the interaction between equity and adoption speed—the pace at which EVs enter the market, measured by their presence in the total vehicle fleet—is therefore central to designing impactful policy.

To explore these dynamics, this study uses Agent-Based Modelling (ABM)—a method that simulates individual-level decisions and emergent social trends—to assess how policy strategies influence EV adoption across socio-economic groups [Ball-Burack et al., 2024]. While past ABM studies have highlighted social contagion and peer influence in the energy transition sector [Varghese et al., 2024], they rarely capture the interaction of evolving barriers and delayed effects over time, and often focus on the broader sector rather than EV-specific dynamics [Jain et al., 2024]. This study addresses that gap by examining the complex interplay of incentives, infrastructure, behavioural tendencies, and demographic diversity across a variation future scenarios. Particular attention is given to targeted policy interventions—measures specifically designed to support under-represented or disadvantaged groups, applied at the optimal time to accelerate market transformation while promoting inclusiveness. Through this lens, the model can be used explore which interventions are most effective in speeding up adoption while ensuring a fair distribution of benefits.

This study is guided by the central research question:
What impact do targeted policy interventions have on the speed and equity of electric vehicle adoption trends in The Hague?

To address this, the study will focus on the following sub-questions:

1. What has been the impact of implemented EV policies on the speed and equity of EV adoption?
2. How do different future scenarios influence the speed and equity of policy interventions?
3. What are the trade-offs between equity and speed objectives in relation to targeted policy interventions to promote EV adoption?

The thesis is structured as follows: Chapter 2 reviews relevant literature and context to answers SRQ1. This is followed by Chapter 3, which details the research design. Chapter 4 then develops this into a conceptual and subsequently a formal approach to the problem. Chapter 5 examines whether the model is fit for purpose, explores its baseline dynamics, and uses these insights to develop the final experimental design. Chapter 6 then systematically presents the results of this design to answer SRQ2 and SRQ3, ultimately addressing the main research question. Chapter 7 offers a critical discussion of the results from academic and societal perspectives, leading to the final conclusions in Chapter 8.

2. Literature and Context

The central aim of this chapter is to answer SQR 1: *What has been the impact of implemented EV policies and initiatives on promoting both equitable and rapid EV adoption?* By drawing on empirical and theoretical insights, the review identifies which interventions have historically shaped adoption trends and where they have fallen short in terms of equity or speed. These findings inform the selection of four policy strategies for further investigation. The chapter concludes by justifying this selection and linking it to the modelling choices in the remainder of the thesis.

2.1. Literature Review

This literature review provides the analytical foundation for the study by addressing three key areas: the determinants of electric vehicle (EV) adoption and their possible policy interventions, the determinant's integration into consumer choice models and frameworks for assessing transport justice. Together, these perspectives clarify how adoption unfolds across diverse contexts and population groups.

2.1.1. Determinants of EV Adoption

Electric vehicle (EV) adoption has become a major interdisciplinary research focus, with the number of publications growing eighteenfold over the past decade [Bhat and Verma, 2022]. As the field matures, scholars have begun to systematize the vast range of influencing factors. Notably, frameworks by Bhat and Verma [2022] and Quaglieri et al. [2024] distinguish between consumer characteristics such as psychological, social, and economic preferences and the technical attributes of the vehicles themselves. Across these frameworks, three dominant domains consistently emerge: cost-related considerations, range-related concerns, and socially embedded behavioural routines.

Range-Related Concerns

Despite advances in battery technology, with current EVs offering an average range of 412 km, range anxiety remains a dominant barrier to adoption. About 77% of Dutch drivers still cite it as a primary concern, despite averaging only 23 km of daily travel [ANWB, 2025]. This discrepancy reveals that range anxiety is more perceptual than logistical in nature.

Empirical evidence shows that range anxiety is linked to risk aversion and increases willingness to pay (WTP) for reliable charging access [Pei et al., 2025]. User satisfaction and perceived ease of use have been shown to significantly reduce range anxiety. Psychological traits such as control beliefs, ambiguity tolerance, and subjective range competence also shape how comfortable drivers feel with their vehicle's range [Franke et al., 2011]. These findings point to the value of interventions like training, experience-based interfaces, and real-time driving feedback as complements to technological improvements.

Infrastructure and geography further modulate range concerns. Wang et al. [2023] find that drivers weigh both distance-related and time-related anxieties when deciding when and where to charge. In regions with poor infrastructure, every 10% increase in remaining battery before the next stop reduces the likelihood of delaying charging by only 9.5%, compared to over 16% in well-equipped regions. These patterns suggest that perceived charging security heavily influences user behaviour. This is further emphasized by current adoption statistics from ANWB [2025] that show that households with a private charger installation options are significantly more likely to buy an EV then those who have not. These insights indicate that for urban EV adoption, where many neighbourhoods lack private charging infrastructure—interventions such as fast charging en route, nearby, or at work could significantly reduce range anxiety.

Planning models reinforce this point (e.g. Sikder et al. [2023]), but also again stress the psychological point, also in connection to grid capacity and budgetary constraints. An optimization model for Amsterdam by Mashhoodi and Van Der Blij [2020] estimated that if drivers use 90% of their battery before charging, only a 31% increase in charging stations is needed to support a sevenfold rise in EVs. However, if drivers charge early at 30% battery, infrastructure needs jump by 167%, with nearly €5 million in added costs. EU policies now increasingly include demand-response measures (e.g., dynamic pricing) and direct charging controls (e.g., minimum battery thresholds) to modify user behaviour rather than expanding infrastructure alone.

Cost Considerations

Pricing remains a critical barrier, despite the total cost of ownership (TCO) often being favourable due to lower fuel, maintenance, and tax expenses [Hopkins et al., 2023]. However, these long-term savings are frequently misunderstood or inaccessible. Siebenhofer et al. [2021] argue that TCO advantages only influence purchasing behaviour when paired with consumer knowledge, which they suggest can be enhanced through transparent and user-friendly financial comparisons.

Nevertheless, the high upfront purchase price of EVs remains a serious obstacle, particularly for lower-income households, who may lack both the capital and credit access to invest in an EV, even if long-term savings are evident. This issue of investment accessibility leads to under-adoption precisely among the groups who stand to benefit most from reduced operational costs. While targeted subsidies are often proposed as a solution, recent research by Joshi et al. [2022] shows that government policy can mediate EV adoption through multiple channels. In their study, price emerges as the most influential factor shaping adoption intentions, but its impact is significantly strengthened when filtered through effectively adapted government actions. Their findings suggest that a broader set of policy instruments should be considered, including financing support, public education campaigns, and strategic promotion initiatives that amplify consumer knowledge while directly addressing affordability barriers.

Fixing these price-related constraints is not only economically important, it also connects to behavioural and social dynamics. A study in Vietnam found that even emotionally satisfied millennial consumers were highly price-sensitive: when price was only altered slightly, the predictive effect of satisfaction on repurchase intention fell by over 13% [Lee et al., 2021]. This demonstrates that even modest price changes can disrupt ingrained consumer preferences and habits, providing a key opportunity for policy to shift behaviour at scale.

Social and Behavioural Factors

As the discussed traditional barriers diminish, social and behavioural factors are projected to grow in influence [Paradies et al., 2023]. Peer effects, social image, and habitual routines all significantly shape EV adoption. Varghese et al. [2024] and Hopkins et al. [2023] show that peer behaviour, especially within local networks, can normalize EV ownership and drive uptake. Social identity also matters: Li et al. [2023] find that perceptions of environmentalism and status play a major role, particularly in urban and affluent settings. These conditions, which promote non-linear growth effects, can be leveraged through policies aimed at helping EV adoption gain a foothold in socio-demographic groups that have so far lagged behind. Song and Potoglou [2020] also identified this effect and began translating the evidence into policy recommendations. They found that carefully designed policies targeting potential consumers with different socio-economic characteristics such as tailored subsidies, local vehicle exchange programs, or promotional campaigns, can be particularly effective

However, social influence is moderated by behavioural inertia. Longitudinal studies by Paradies et al. [2023] and Zhang et al. [2024] show that roughly 30% of consumers exhibit "routine" buying behaviour, preferring vehicles similar to previous choices. Unless disrupted by strong external stimuli, these consumers tend to exclude EVs from their consideration set. As mentioned earlier, financial incentives or clear value propositions can push even routine buyers toward EVs. However, this effect is relatively small and primarily visible among younger generations. Research into repurchasing behaviour—such as the study by Lee [2020] on mobile phones—shows that approximately 70% of repurchase decisions stem from consumer satisfaction. Studies like these have led to the insight that though stimulating EV adoption may not be sufficient to counter repurchasing behaviour; much can also be gained by making non-EVs less attractive. This idea has gained traction in recent years and has contributed to the implementation of zero-emission zone policies in the centres of many major EU cities [Gao et al., 2022]. In the Netherlands, these zones currently only exclude high-emission diesel vehicles, but there is a strong argument for extending the restrictions to all internal combustion engine vehicles, regardless of fuel type [Gemeente Den Haag, 2025].

Notably, once a consumer does switch to EV it tends to stick: [Hertzke et al., 2025] reports that although EV buyers often switch brands (up to 66% in China), only 1% show interest in returning to internal combustion engine vehicles. This suggests that once users adapt, retention is high.

Technological and Economic Advancement

Technological innovation and economic development are central to the future of EV adoption, interacting with behavioural, infrastructural, and spatial determinants. Shared mobility platforms, particularly those offering electric fleets, continue to provide low-risk, low-cost access to EVs. Around 20% of users would forgo a second car if car-sharing were available nearby [Liao et al., 2018]. This effect is especially pronounced in dense urban areas, where high ownership costs and multimodal transport systems create ideal conditions for shared mobility. Moreover, regular exposure to EVs through such services helps reduce psychological barriers like range anxiety [Li et al., 2023].

Urban form strongly influences these dynamics. Cities with compact layouts and integrated transit systems not only facilitate EV infrastructure deployment but also reinforce social norms around sustainable mobility [Mulley et al., 2020]. This creates feedback loops in which infrastructure visibility supports normalization and adoption.

Technological advancements such as improvements in charging speed, battery range, and vehicle-to-grid integration continue to lower adoption barriers. However, benefits are unevenly distributed. In less urbanized or lower-income regions, infrastructure expansion may lag, creating access disparities [Gnann et al., 2023]. Economic conditions further shape adoption: while high-income consumers are driven by values and convenience, lower-income groups remain price-sensitive and face access limitations despite favourable long-term costs [McDonald et al., 2022]. Without targeted subsidies or financing mechanisms, adoption in these groups is likely to stagnate.

Looking ahead, EV adoption is expected to thrive in high-tech, urbanized contexts with strong infrastructure and inclusive policies. Conversely, futures marked by unequal growth or weak public investment risk amplifying disparities in access. Pasaoglu et al. [2015] argues that therefore scenario-based planning is essential to ensure that technological progress aligns with behavioural readiness and spatial equity. Flexible, data-driven policies such as regionalized incentives, local infrastructure planning, and targeted outreach can help ensure scalable and inclusive transitions.

2.1.2. Consumer Decision Process

As the economic, psychological, and social determinants of electric vehicle (EV) adoption become better understood, the challenge shifts toward accurately modelling how these factors interact in actual consumer decisions. This requires not only acknowledging individual heterogeneity and bounded rationality but also understanding how perceptions, context, and policy signals shape choices across time. Simple projections based on price elasticity or total cost of ownership are insufficient, especially when emotional, social, and habitual factors dominate decision-making for a significant share of the population [Quaglieri et al., 2024; Bhat and Verma, 2022]. Therefore, the question is no longer only what drives EV adoption, but how these drivers play out dynamically within the decision-making process of real consumers.

Consumer Decision Theories

To select a behavioural decision model for simulating household EV adoption, several frameworks were considered, each with varying strengths in capturing dynamic, phase-based, and policy-relevant behaviour.

One well-known approach is the Theory of Planned Behaviour (TPB), which models intention formation based on attitudes, social norms, and perceived control. TPB has been applied to EV adoption [Buhmann et al., 2024], but its static and linear structure limits its ability to capture feedback, evolving contexts, or repeated exposures over time. This makes it less suitable for simulations that aim to reflect changing behavioural processes.

Discrete choice models, such as multinomial or mixed logit, offer a structured way to represent utility-based decisions by incorporating costs, vehicle range, and policy incentives. These models are common in transport studies, but often assume stable preferences and fully rational decision-making. As a result, they do not easily account for bounded rationality, social influence, or intention formation that evolves through experience or exposure [Rudolph, 2016].

Other approaches, including the Value-Belief-Norm theory and the Norm Activation Model, emphasize moral motivations, pro-environmental identity, or normative beliefs. While these frameworks offer greater psychological depth, they often lack clear mechanisms for modelling how intentions form or respond to policy within a simulation context. Similarly, interpretive and qualitative methods, such as discourse analysis, have been used to explore how consumers frame emerging technologies like EVs [Jain et al., 2024]. These approaches provide valuable insight into meaning-making but are difficult to scale or incorporate into simulation-based policy analysis.

To address these limitations, the CODEC model developed by Paradies et al. [2023] offers a simulation-ready framework that combines key strengths of the previously discussed approaches. Like discrete choice models, it can represent utility-based decision factors such as cost, range, and incentives, but it also incorporates the evolving, phase-based processes absent in static frameworks like TPB. Similar to value-based theories, CODEC can account for psychological drivers such as environmental concern or social influence, yet it does so in a form that is modular and quantifiable for simulation.

In CODEC, consumer behaviour is modelled as a sequence of three phases: attention, enablement, and intention. This structure allows for dynamic change over time as households respond to shifting infrastructure, policy measures, or social dynamics. By explicitly modelling how policy interventions affect each phase—for example, increasing visibility of incentives (attention), improving perceived accessibility (enablement), or strengthening willingness to adopt (intention)—CODEC can capture both behavioural depth and simulation compatibility. This makes it well suited for representing heterogeneous, adaptive EV adoption pathways in response to multiple, interacting policy levers [TNO, Sustainable Urban Mobility Safety, 2023].

In CODEC, it is possible to simulate how different policy interventions affect attention (for example, visibility of incentives), reshape enablement (such as perceived accessibility), and influence intention (such as willingness to adopt) over time. Compared to TPB, discrete choice models, or value-based frameworks, CODEC provides a more behaviourally grounded and simulation-compatible structure for modeling consumer transitions in EV adoption [TNO, Sustainable Urban Mobility Safety, 2023].

Modelling approaches

To operationalise the CODEC decision framework, several modelling paradigms were assessed using the decision diagram from Novizayanti et al. [2021] (Figure 2.1), which compares features such as feedback mechanisms, system evolution, autonomy, heterogeneity, and explicitness of decision-making.

First, the model must capture dynamic feedback loops, including peer influence, infrastructure responsiveness, and policy adaptation over time. This rules out equation-based statistical models, which assume fixed relationships and lack temporal dynamics.

Second, as the system evolves due to technological change, shifting policies, and changing social norms, system dynamics (SD) modelling offers strengths in analysing feedback and macro-level patterns. However, SD is less suited to representing heterogeneous household decisions and becomes unwieldy when feedback structures change over time. Hybrid SD–discrete choice models [Huang et al., 2021] improve behavioural realism but often have limited scalability and transparency.

Third, evolutionary modelling can simulate adaptation through selection and variation but lacks goal-directed agents; change emerges at the population level rather than from intentional, phased decisions as in CODEC, making it unsuitable here.

Fourth, the system involves heterogeneous households with diverse incomes, car ownership histories, travel needs, and social networks, interacting with evolving infrastructure and policy contexts. This complexity rules out cellular automata (CA), which simulate spatial diffusion using uniform rules but lack explicit, context-dependent decision-making.

Given these considerations, agent-based modelling (ABM) is the most suitable approach. ABMs simulate interactions between diverse agents, policies, and infrastructure in dynamic settings. As described by Nikolic and Ghorbani [2011], this involves five key phases: system analysis, model design, detailed design, implementation, and evaluation.

Integrating the CODEC decision framework into the ABM addresses limitations in existing EV adoption models, which often under-represent psychological variation [Ball-Burack et al., 2024], poorly integrate decision context and policy timing [Novizayanti et al., 2021], and overlook household routines and mobility practices [Querini and Benetto, 2014].

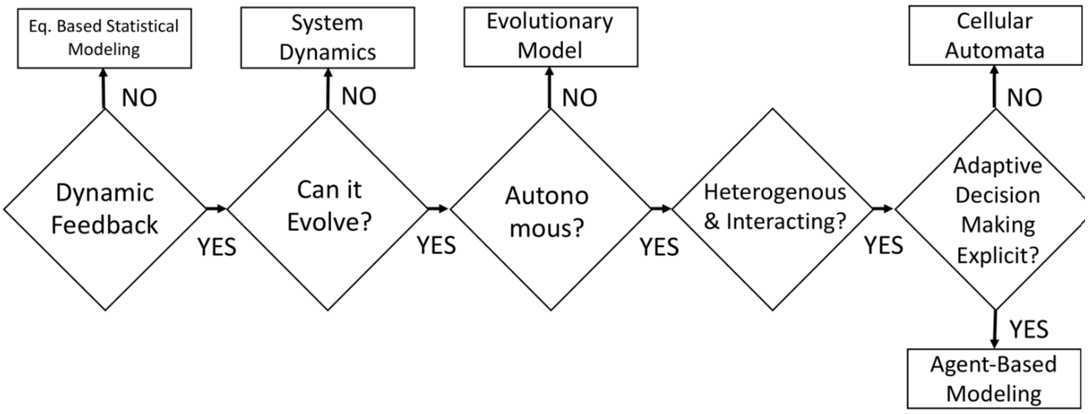

Figure 2.1.: Decision diagram of the modeling method from Novizayanti et al. [2021].

2.1.3. Equity in EV Policy Measures

The transition to electric vehicles (EVs) brings forward fundamental questions of distributive justice and calls for a deeper understanding of what constitutes fair access to mobility, infrastructure, and public support. As Mandolakani and Singleton [2024] argue, transport equity is not a singular objective but a contested concept involving trade-offs across spatial, social, and intergenerational dimensions. Policymakers are thus confronted with the task of deciding which dimensions to prioritise and how to do so in ways that are effective, measurable, and transparent.

Equity and Justice Theory

A useful distinction is between *horizontal equity*, which treats groups with similar needs in the same way, and *vertical equity*, which justifies allocating more resources to disadvantaged groups in order to address unequal starting points. The *ladder of justice* [Martens and Golub, 2018] offers a further lens for assessing policy ambition. At its lowest level, explicit nondiscrimination ensures formal equality but ignores disparate outcomes. Pareto improvement requires that no group be made worse off, even if benefits are uneven. Proportional equity seeks similar benefit levels across groups. At the highest rung, restorative justice actively allocates greater benefits to groups that have been disadvantaged in the past [Hopkins et al., 2023]. Given that higher-income households have so far captured most EV subsidies, restorative strategies are particularly relevant, although proportional or Pareto-based approaches may still appeal to those favouring balanced or incremental gains.

Political philosophy provides additional perspectives. Libertarianism emphasises individual choice and minimal intervention, often tolerating unequal outcomes if they result from voluntary exchange. Utilitarianism prioritises maximising overall welfare, even if gains are unevenly distributed. Egalitarian approaches, such as Rawls' difference principle, seek to improve the position of the least advantaged, while sufficientarianism focuses on ensuring everyone meets a basic threshold of access. Amartya Sen and Martha Nussbaum's capabilities approach frames mobility as a fundamental capability, highlighting the importance of both opportunity and the means to participate fully in society.

Dimensions and Measurement

Within this conceptual landscape priorities in general two directions are dimensions are accepted. The first is spatial equity and can be assessed by analysing the geographical distribution of charging infrastructure. Studies by Sikder et al. [2023] and Soltani Mandolakani and Singleton [2024] show that EV chargers are often concentrated in wealthier, urban areas with better grid access and private parking. Metrics here might include charging density per capita or per registered vehicle across income groups or postal code areas, revealing whether rural or lower-income districts are underserved.

The second dimension is social equity and concerns focus on affordability and demographic disparities in EV ownership. While EVs may reduce operating costs over time, high upfront prices and rebate structures often exclude lower-income households. Research by Soltani Mandolakani and Singleton [2024] and Hopkins et al. [2023] finds that subsidies have disproportionately benefited wealthier consumers, a regressive outcome. Measuring this involves tracking the socio-economic profile of EV owners or subsidy recipients and comparing it to the general population.

Addressing these inequities requires targeted interventions. Hopkins et al. [2023] highlight measures such as subsidies for low-income households, public overnight charging in dense neighbourhoods, and community-led outreach. Expanding access to second-hand EVs, battery leasing, and shared electric mobility can further reduce barriers. Their effectiveness can be monitored through uptake rates in target groups, modal share changes, or participation in policy schemes.

To determine which interventions are most effective and for whom, these qualitative insights must be translated into quantitative, equity-sensitive modelling. Incorporating demographic variables into consumer agent profiles, simulating differential access to infrastructure, and mapping policy reach across socio-economic groups allows scenario analyses to explicitly test whether vulnerable populations are included or excluded under different transition pathways. Equity, then, must not be treated as a secondary consideration, but as a central design criterion embedded into both the empirical evaluation of current programs and the simulation of future outcomes.

2.2. Context

2.2.1. EV Adoption in Europe and Beyond

The share of electric vehicles (EVs) in new car sales has risen sharply worldwide over the past decade, yet progress varies considerably between regions (see Figure 2.2). In the European Union (EU), EVs accounted for just under 20% of new passenger car registrations in 2023, up from less than 3% in 2019 [Ritchie, 2024]. The United States lags further behind, with EVs representing only about 8% of new sales. In stark contrast, Norway has emerged as the undisputed global leader, with EVs consistently exceeding 80% of new car sales since 2022 (Figure 2.2). Despite this rapid growth, the EU as a whole trails Norway and other early leaders in fleet penetration—the share of total cars on the road that are electric. This metric matters because fleet turnover is slow, Norway is still only at 29%, and climate and air-quality targets depend on replacing the existing internal combustion fleet.

Norway's success stems from a sustained and coherent policy package maintained over more than a decade. Key measures include comprehensive fiscal incentives such as full exemptions from VAT and purchase/import taxes on EVs; cost advantages in use through reduced road tolls, ferry fees, and parking charges; guaranteed access to charging via the "right to charge" law; and strong public engagement by the Norwegian EV Association. Importantly, these measures were implemented early, long before EVs became mainstream, creating cumulative advantages that compounded over time.

2.2.2. Policy Approaches and Effectiveness

Across Europe, countries have experimented with a variety of incentives and regulations to stimulate EV adoption. Data from the European Commission [2024] shows that most EU states rely heavily on purchase subsidies for new EVs, though some, like the Netherlands, have focused on road tax exemptions or support for used EV purchases. The impact of these

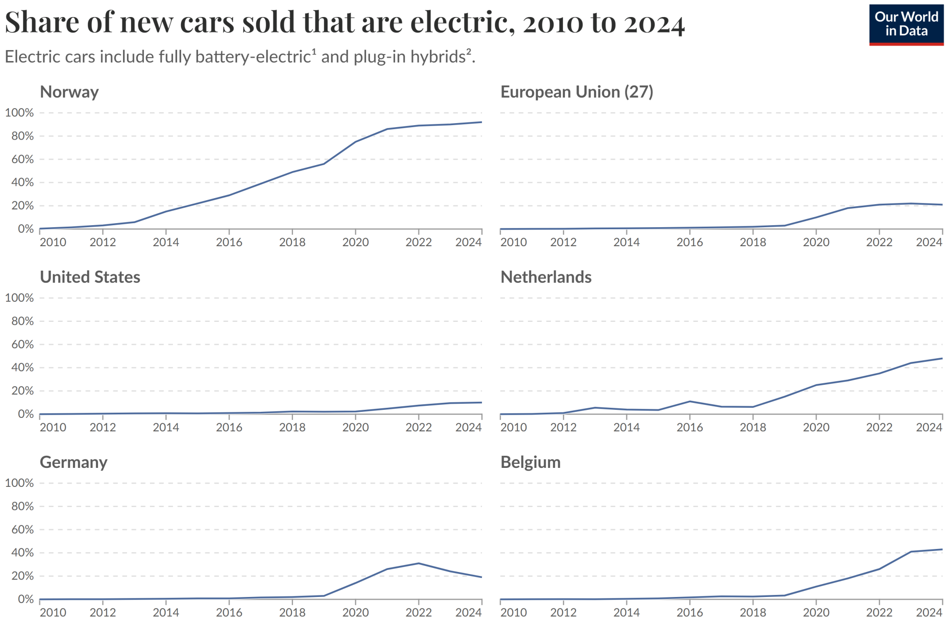

Figure 2.2.: Share of new cars sold that are electric, 2010–2024 [Ritchie, 2024].

measures varies: high adoption rates are associated with countries that combine purchase incentives, strong charging infrastructure rollout, and consistent policy signals. In contrast, sudden subsidy reductions have slowed uptake which you can already see in the end of the figures in Figure 2.2 for Germany and the European Union.

While financial incentives have boosted sales, they have also amplified equity concerns. In the Netherlands, EV ownership is concentrated among higher-income households (see figure Figure 2.3), who are more likely to have off-street parking and the capital to purchase new vehicles. Lower-income groups, often renters in dense urban areas, face persistent barriers, even when subsidies are formally available but simply won't be sufficient to change minds..

With the forecasts of Paradies et al. [2023] suggesting that the will stagnate as well a bit of nothing unannounced will be done. It becomes interesting to see how the government can stop stagnation whilst seriously take into account the apparent inequity.

2.2.3. The Hague as a Case Study

The Hague, located in the densely populated Randstad, offers a telling microcosm for studying how national EV policy interacts with local infrastructure and socio-economic diversity. The Netherlands is among Europe's frontrunners in EV adoption, supported by a dense public charging network and generous fiscal incentives. Yet, as Figure 6.9 shows, adoption within The Hague is far from uniform.

Wealthier, low-density districts—often with detached housing and off-street parking—show EV penetration rates above 10%, while dense, lower-income neighbourhoods remain be-

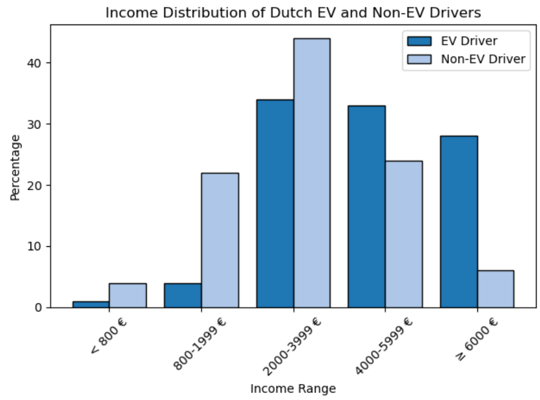

Figure 2.3.: Income distribution of Dutch EV and non-EV drivers. Higher-income households are overrepresented among EV owners. Based on data of European Commission [2024]

low 2%. This spatial divide closely mirrors the distribution of standardised disposable income, underscoring the link between household affluence and EV uptake [Municipality The Hague, 2024]. The Hague is therefore a particularly interesting case, as this effect is especially visible in a city known for its strong contrasts between neighbourhoods. Furthermore, the municipality is a frontrunner in tackling emissions, for instance through participation in national zero-emission zone mobility programmes, and it transparently publishes data and insights on the impacts of these measures.

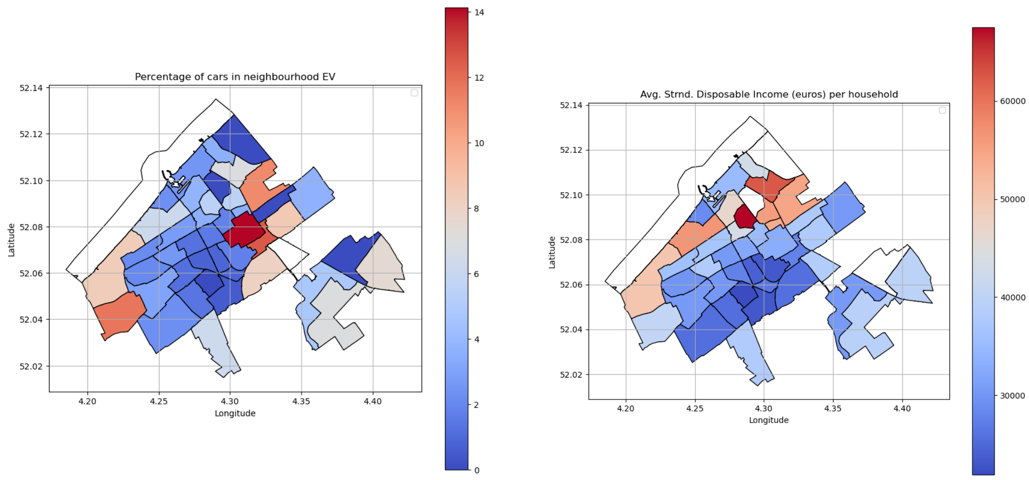

(a)

(b)

Figure 2.4.: (a) Spread of EV adoption levels in The Hague, 2024; (b) Spread of standardised income across The Hague, 2021 (both based on data from [Municipality The Hague, 2024]).

2.3. Conclusion

This literature review has addressed the first sub-question of this study:

> *What has been the impact of implemented EV policies and initiatives on promoting both equitable and rapid EV adoption?*

The evidence shows that while current EV policies have contributed to accelerating adoption, their benefits are unevenly distributed. Higher-income households and well-connected urban areas have captured the largest share of gains, while lower-income and infrastructure-poor neighbourhoods remain underserved. Traditional policy approaches often neglect the behavioural and psychological drivers of adoption, including affordability perceptions, range confidence, and the influence of social norms. Across the literature, three consistent determinants emerge as critical to both speed and fairness: the cost barrier, access to reliable charging, and the visibility and social acceptability of EVs. These determinants are strongly shaped by local context and by how policies interact with existing socio-economic inequalities. Importantly, research suggests that rapid and equitable adoption is achievable if policies are designed to address these barriers in an integrated and targeted way.

Building on these insights, this study identifies four complementary policy interventions designed to address the key adoption determinants while maximising both performance and equity: (1) *targeted subsidies* to improve affordability without regressive effects; (2) *targeted public charging expansion* to reduce range anxiety and close infrastructure gaps; (3) *zero-emission zones* to influence repurchase behaviour and accelerate normative change; and (4) *community-based information campaigns* to improve EV knowledge, trust, and visibility, particularly among hesitant or disadvantaged groups. These measures span financial, infrastructural, regulatory, and informational levers, allowing for diverse and potentially reinforcing effects on household decision-making.

To assess their effectiveness, the next stage of this research will implement these interventions—individually and in combination—within an agent-based model of The Hague, using the CODEC decision framework. This approach allows for the simulation of heterogeneous, phase-based decision-making processes that respond dynamically to evolving infrastructure, policy measures, and social influence. Policy performance will be evaluated using a balanced set of metrics that incorporate both adoption speed and equity from the outset: the percentage of adoption in the lowest-income group, two times the standard deviation of adoption across household archetypes, the time taken to reach 25% adoption, and the total fleet share of EVs by 2035. Together, these metrics will provide a robust basis for comparing trade-offs and identifying policy packages that deliver a fast *and* fair transition to electric mobility.

3. Methods

3.1. Overall Research Design

This thesis explores how electric vehicle (EV) policies can be designed to be both effective and equitable in a context of deep uncertainty. To address this, it employs an innovative research design that integrates a custom-built agent-based model (ABM) with exploratory policy analysis tools and spatial generalization techniques. The approach blends behavioural realism at the micro-level with broad applicability at the macro-level, allowing for both rigorous simulation and practical relevance.

At the heart of the study is an agent-based simulation of household EV adoption behaviour in The Hague. Households are modelled as individual decision-making agents, whose choices are shaped by economic factors, behavioural preferences, social influence, and government policy. To ensure the model captures the diversity of real-world urban environments, a typology of four neighbourhood archetypes was developed using Dutch statistical and spatial datasets. These archetypes—representing combinations of income, infrastructure, and education—form the backbone of a scalable modelling structure that reflects both demographic variety and spatial inequality.

Policies are not tested in isolation, but as policy packages, acknowledging the complex interactions between infrastructure, subsidies, regulation, and awareness campaigns. The Exploratory Modelling and Analysis (EMA) Workbench is first used to design and evaluate a stratified sample of over 200 policy-scenario combinations. Then based on performance and diversity five are experimented with further, by analysing their performance in 1000 different scenarios. This allows for robust testing of interventions across a wide range of plausible futures, identifying not only which policies perform best, but under what conditions and for whom.

Finally, the analysis framework includes a multi-objective evaluation of performance. Adoption speed, final EV share, distributional equity, and inclusion of low-income groups are assessed and compared using both scenario-dependent metrics and aggregate trade-off tools such as Pareto frontiers and parallel coordinate plots. This dual focus—on both performance and fairness—ensures that recommended strategies align with the broader societal goals of a just and sustainable transition.

Taken together, this research design offers a novel contribution by combining spatial abstraction, behavioural modelling, and deep policy experimentation into one coherent and replicable framework.

3.2. Methodological Approach

3.2.1. ABM Modelling Approach

Taking note of the structure laid out by Nikolic and Ghorbani [2011], the methods for creating and experimenting with the ABM follow a structured approach:

1. *System Analysis*, detailed in Chapter 2, includes both the problem and system identification. It frames the societal challenge of electric vehicle (EV) adoption and defines equity-oriented policy objectives. It also identifies relevant behavioural and spatial system elements, which inform the scope and structure of the model.

2. *Conceptualisation and Formalisation*, as described in Chapter 4, involves translating the problem and system understanding into a formal agent-based structure. Households are modelled as autonomous agents with socio-demographic attributes, behavioural decision-making logic, and social interaction mechanisms. Policy instruments, such as subsidies, infrastructure expansion, and awareness campaigns, are integrated as time-dependent external influences.

3. *Implementation and Verification*, described in Chapter 5, includes model coding and iterative debugging. The model is calibrated against historical EV adoption data (2020–2025), and a sensitivity analysis is performed to identify and refine key parameters. Model behaviour is assessed against real-world plausibility and known patterns.

4. *Simulation and Analyses*, detailed in Chapter 6, consists of extensive policy experimentation under uncertainty. This step leverages the EMA Workbench to explore 1000 scenario-policy combinations and evaluate each policy's robustness and trade-offs impact. This supports strategic policy selection.

Although this outline presents the steps linearly, in practice they involved frequent iteration and refinement. Once model performance was deemed sufficient to answer the research questions, analysis transitioned to the evaluation and reflection chapter 7.

3.2.2. Archetypical Approach

A key methodological innovation of this study is the use of spatial-socioeconomic neighbourhood archetypes to combine high-resolution household-level modelling with broader generalizability. Rather than modelling each individual district in The Hague, four archetypes were identified based on clustering demographic and structural characteristics such as income, density, and education level. These archetypes serve as proxies for real neighbourhoods, allowing the model to simulate the effects of policy interventions on diverse but representative community types. This structure enables scalability, realism, and tractability, supporting both detailed behavioural insight and policy-level relevance.

3.2.3. Data Sources

Reliable and context-specific data is essential for the validity of agent-based models. This study makes extensive use of open public datasets, primarily from the Centraal Bureau voor de Statistiek (CBS) and the 'Den Haag in Cijfers' portal. These sources provide detailed socio-demographic, spatial, and economic information at the neighbourhood level, which was used to instantiate and cluster household agents. Additional insights into consumer behaviour, EV adoption trends, and policy impacts were drawn from the academic literature and expert consultations at TNO. Combining these data sources allows the model to reflect both statistical realism and contextual nuance.

3.2.4. Research Scope

The focus of this research is the urban setting, for two main reasons. First, the environmental consequences of non-electric vehicle use are more immediately pronounced in cities, as exhaust gases tend to accumulate in the air and cannot easily disperse from densely built areas. Research has shown that this health hazard is particularly severe in lower-income neighbourhoods, where people are more likely to drive older, more polluting cars [Hopkins et al., 2023]. Second, the urban context is particularly interesting because infrastructure availability, spatial constraints, and social dynamics strongly influence EV adoption patterns. Different outcomes in terms of equity versus performance can be expected, making it a relevant setting for this research to contribute to the literature by developing an innovative way to incorporate these different aspects and explore needs under diverse future scenarios, particularly in light of the fast-changing conditions in urban mobility. For example, the research includes scenarios in which *No Car* ownership trends grow significantly, possibly due to the availability of far superior public transport.

In contrast, suburban and rural settings exhibit different behavioural and infrastructural mechanisms due to differing mobility needs and lower housing densities; therefore, they fall outside the scope of this analysis.

The case study focuses on The Hague, a city characterized by significant socio-economic diversity, a compact urban form, and active zero-emission zone initiatives. Its rich availability of open-access data and detailed municipal statistics provides a robust empirical foundation for modelling. At the same time, its structural features make it a relevant analogue for other mid-sized European cities facing similar mobility transitions.

4. Model Conceptualization and Formalization

4.1. Conceptual Framework

4.1.1. Model Purpose

The agent-based model (ABM) simulates household-level electric vehicle (EV) adoption over time, aiming to explore how behavioural, economic, and infrastructural factors interact under different policy scenarios. The model captures both structural and behavioural drivers of car ownership decisions, with a particular focus on the transition from conventional to electric vehicles. It integrates the CODEC framework comprising Attention, Enablement, and Intention phases to represent the cognitive and contextual processes underlying household decision-making. In doing so, the model allows for differentiated responses across household types and helps assess the effectiveness and equity of targeted interventions.

4.1.2. Time and Space

The model operates on an annual time step, representing evolving adoption dynamics over a 14 year period of which the first 3 are already in the past. Technological progress, policy changes, and social diffusion are implemented incrementally, reflecting realistic temporal developments such as improving EV ranges, declining costs, and maturing second-hand markets.

The environment is spatially explicit and grid-based, with households distributed using income-based clustering (via KMeans). This setup mirrors spatial segregation in real-world neighbourhoods, where residents with similar socioeconomic characteristics often live in proximity. This spatial dimension enables analysis of localized diffusion patterns and social influence within peer groups.

4.1.3. Overview

Figure 4.1 presents a high-level overview of the conceptual ABM design, adapted from Huang et al. [2021]. The design consists of household agents making vehicle purchase decisions, a government agent implementing policy interventions, and a system environment that shapes the context perceived by the agents.

Each household represents a potential car-owning entity with heterogeneous attributes, such as income (operationalized through car budget), daily driving needs, EV knowledge, and car

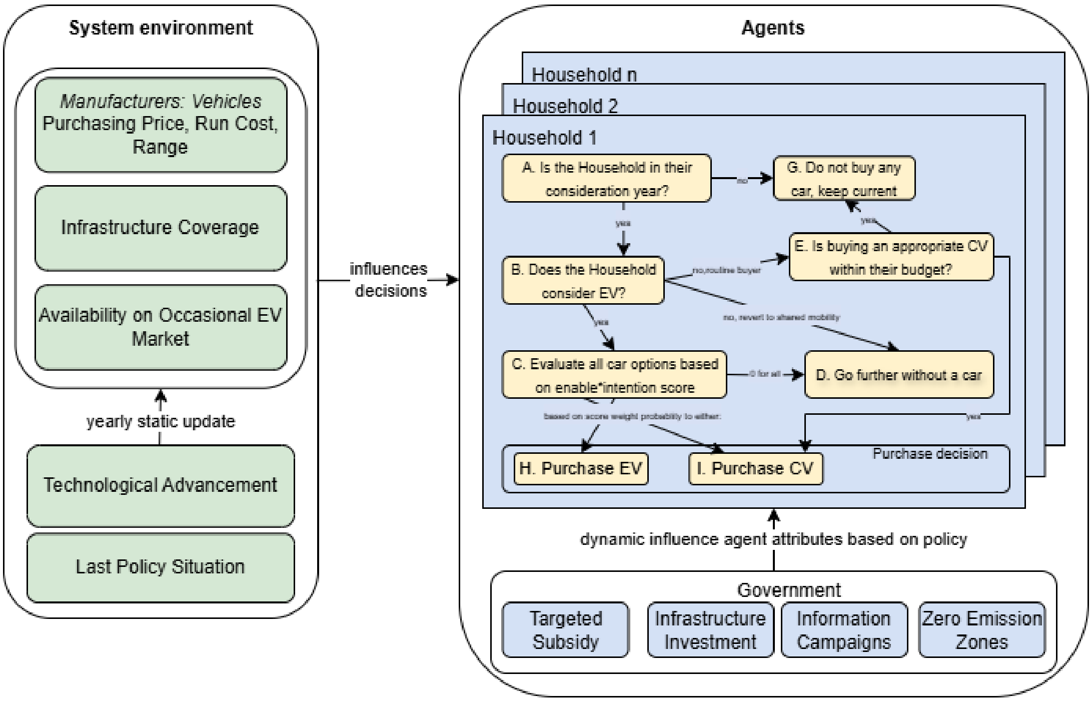

Figure 4.1.: Overview of conceptual ABM setup, adapted from Huang et al. [2021].

ownership preferences. Some households exhibit a "no-car preference" for lifestyle or environmental reasons, while others are modelled as "routine buyers" who repeat past purchase behaviour unless strongly triggered to reconsider.

Households periodically enter a decision cycle, typically triggered when a vehicle reaches its replacement age. However if the agent is a routine buyer, no deliberation occurs, instead the agent immediately decides to purchase a conventional vehicle again. If the decision cycle is entered, it follows the structure of the CODEC framework: in the Attention phase, the household becomes aware of the need to make a decision; in the Enablement phase, it evaluates the feasibility of different options based on factors such as EV knowledge, budget, infrastructure availability, range adequacy, and vehicle availability; and in the Intention phase, motivational aspects like running costs, product features, and social influence are considered

Social influence is modelled using a Moore neighbourhood structure to simulate peer effects within the same street or local environment, reflecting the social comparison theory of Festinger [1954]. The determinants of agent behaviour and their implementation across CODEC phases are visualized in Figure 4.2.

Price sensitivity, range anxiety, social influence, and habitual behaviour are operationalized through respective CODEC mechanisms. Car purchase decisions result in a probabilistic selection between four vehicle types (new/used, EV/CV), or opting out of ownership altogether, based on weighted enablement and intention scores.

Technological and regulatory changes are incorporated incrementally: EV range and affordability improve annually, second-hand EVs become more available, and tax exemptions are

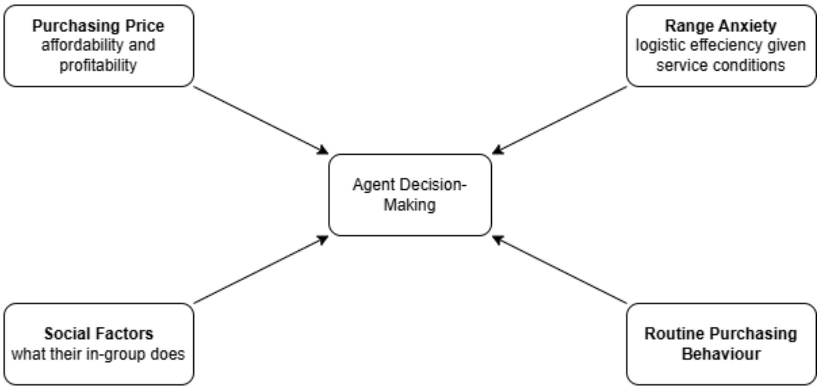

Figure 4.2.: Determinants of agent behaviour in EV decision-making.

gradually phased out. A government agent can implement targeted interventions—such as subsidies, infrastructure investment, public awareness campaigns, and low-emission mandates at predefined time points, constrained to biannual updates between 2026 and 2030.

Key model outputs include the shares of EV-owning, CV-owning, and car-free households. These outputs can be further operationalised to evaluate the speed and equity of EV adoption when compared across time and different neighbourhood types.

4.2. Formal Model

This section formalises the agent-based model simulating household-level electric vehicle (EV) adoption. It outlines the system structure and dynamics, details agent decision-making logic, and specifies policy implementation. To support understanding, Figure 4.3 is presented here to provide a clear overview of the operational flow between the various elements discussed throughout the remainder of the section.

Simulation Pseudocode

```
FOR each time step t in simulation:

    // 1. Apply time-specific government policies
    IF t in policy_schedule:
        Undo previous policies
        Apply current policies

    // 2. Each household agent takes a step
    FOR each household agent:
        Increase car age by 1

        IF car is due for replacement:
            IF agent is a routine buyer:
                Repurchase same vehicle type
            ELSE:
                IF agent has no-car preference or opts for car-free:
                    Update ownership to "no car"
                ELSE:
                    FOR each available vehicle type:
                        Compute enablement score (feasibility)
                        Compute intention score (motivation)
                        Multiply both for total score
                    IF any scores > 0:
                        Select vehicle probabilistically based on scores
                        Update ownership
                    ELSE:
                        Possibly switch to car-free

    // 3. Apply annual technology changes
    Increase EV infrastructure coverage (bounded to 100%)
    Decrease EV purchase price
    Increase EV range
    Adjust EV running costs (policy-dependent)
    Increase second-hand EV availability

    // 4. Collect model-level statistics
    Record shares of EVs, CVs, non-car households, EV fleet, LEHs
```

Figure 4.3.: Pseudocode overview of the simulation model's operational flow.

4.2.1. System Initialization and Dynamics

The agent-based model is implemented in Python using the `Mesa` framework and simulates the adoption of electric vehicles (EVs) over time. The model is encapsulated in the class `EVAdoptionModel`, which governs the initialization, spatial structure, agent creation, and time-step evolution of the system. Table 4.1 indicates the main parameters of the system.

Table 4.1.: Overview of System Initialization Parameters

Parameter	Value (default (2022) → range)	Description
Global Parameters		
`n`	1000 → int >0	Number of household agents; can be adjusted to the actual number of households per neighbourhood based on [Central Bureau Statistics Netherlands CBS, 2023].
`stepn`	0 → int	Model step counter.
`perc_EV`	0.03 → [0, 1]	Initial share of EV-owning households in a neighbourhood; default reflects the Dutch average [Central Bureau Statistics Netherlands CBS, 2023].
`perc_NC`	0.26 → [0, 1]	Share of non-car-owning households; default based on the Dutch average in 2022 [Central Bureau Statistics Netherlands CBS, 2023].
`grid`	MultiGrid(torus=True)	Spatial model grid with wrapped edges.
`ev_infra_coverage`	0.615 → [0, 1]	Perceived EV infrastructure sufficiency; default from [Paradies et al., 2023].
`ev2avail`	0.3 → [0, 1]	Availability of second-hand EVs [ANWB, 2025].
`clustering`	False → {True, False}	Enables or disables spatial clustering of households.
Vehicle Options (Price, Cost, Range)		
`EV_h`	€42k, €145, 343km	New electric vehicle, specifications from [ANWB, 2025].
`EV_l`	€20k, €174, 309km	Second-hand (occasional) electric vehicle [ANWB, 2025].
`CV_h`	€35.5k, €268, 750km	New combustion vehicle [ANWB, 2025].
`CV_l`	€12k, €322, 750km	Second-hand (occasional) combustion vehicle [ANWB, 2025].

System Initialization

At initialization, the model sets up a population of n agents (`n`), with initial electric vehicle (EV) ownership and non-car households defined by `perc_EV` and `perc_NC`, respectively. The spatial environment is configured as a two-dimensional toroidal grid (`grid`) that fits the population as tightly as possible. Agents are placed spatially clustered (`clustering`) using k-means based on income, with three clusters selected argued by the elbow method. In clustered setups, agents with similar income levels are co-located to reflect neighbourhood-level socioeconomic stratification. Note that due to rounding in grid dimensions, some blank cells may remain; these represent corner houses with fewer immediate social neighbours.

The model includes vehicle manufacturers offering both electric and combustion vehicle (CV) options, differentiated by segment—e.g., `EV_l`, `EV_h`, `CV_l`, and `CV_h`—with specific prices, ranges, and running costs. EV infrastructure availability is captured by `ev_infra_coverage` and can increase over time. The second-hand EV market is also represented (`ev2avail`), influencing affordability and access.

Additional to the main parameters indicated in the table, a comprehensive set of behavioural scaling parameters (`d_..`, range [0.5,2]) is included in the model to enable sensitivity analyses on structural and behavioural dynamics, including budget and price distributions, EV awareness, routine purchasing, social influence, infrastructure variation, and technological evolution (e.g., `d_tech_price`, `d_tech_range`). These allow for robust exploration of how assumptions impact both the speed and equity of EV adoption.

System Advancement

At each discrete time step, the model progresses through a structured update cycle to reflect behavioural, technological, market, and policy changes. First, all agents execute their individual decision routines via the `step()` method. During this phase, households may decide whether to purchase a new vehicle, influenced by economic conditions, social network dynamics, and local infrastructural availability (see Section 4.2.2).

Technological development occurs annually. Specifically, EV infrastructure coverage expands incrementally by $0.041 \times$ `d_infra_cov`, up to a maximum of 100% . Simultaneously, technological advancements in EVs are modelled by applying a 5% annual reduction in vehicle prices, scaled by the `d_tech_price` parameter, and also a 5% increase in driving range, scaled by `d_tech_range`. To simulate policy shifts, EV running costs also increase by €15 in year 3 and €30 in year 4, reflecting the gradual phase-out of road tax exemptions. These numbers are all economic and technological projections following from data of [Paradies et al., 2023] and [ANWB, 2025].

The model accounts for changing market dynamics by increasing the availability of second-hand EVs each year [Hopkins et al., 2023]. This growth is modelled as a multiplicative increase by an factor of $1.2 \times$ `d_oc_avail`, capturing the expanding influence of the used EV market over time.

Policy implementation is handled dynamically through the `timed_policy_configs` dictionary. This allows the model's `Government` object to apply or reverse specific interventions at predefined time steps, enabling simulations of changing regulatory landscapes (see Section 4.2.3) .

Finally, key system metrics are collected at every step using `mesa.DataCollector`. These include the proportion of EVs, combustion vehicles, and non-car households, as well as the composition of the EV fleet and the percentage of zero-emission households.

4.2.2. Agents: Household Initialization and behaviour

Each agent in the model represents a household making structured vehicle ownership decisions. The decision-making logic is based on a hybrid utility and behaviour-theoretic framework, based on the CODEC model. Table 4.2 indicates the main parameters of the agents.

Household Initialization

At model startup, each Household agent is initialized with heterogeneous socio-economic and behavioural attributes. This heterogeneity is produced by sampling each agent's attributes around the known average, with a reasonable standard deviation and probabilistically setting their initial state variables such as verhicle ownership state.

Vehicle ownership status is assigned probabilistically based on the system parameters `perc_EV` and `perc_NC`, distinguishing combustion vehicle (CV) owners, electric vehicle (EV) owners, and non-car households. Among the latter, 81% (texttno_car_preferrance) are set to express a strong intrinsic preference for not owning a car and will therefore not consider one again. Each household's actual daily travel distance (`daily_km`) is sampled from a truncated normal distribution with a minimum of 5 km. This informs their daily range requirement, which is different from the perceived range variable (`range_need`), sampled similarly from a normal distribution but typically around 20 times larger

Financial constraints on vehicle purchasing are captured by sampling budgets for high- and low-tier EVs and CVs from normal distributions centred on empirical averages, each with a 25% standard deviation. Similarly, the maximum acceptable running cost (`max_running_cost`) is drawn from a scaled normal distribution to reflect variation in cost sensitivity. Behavioural traits include an initial EV knowledge level, as well as a probabilistic assignment of a routine buyer flag to model habitual vehicle replacement patterns. Finally, each household's vehicle replacement age is drawn from a log-normal distribution, and the current vehicle age is initialized uniformly within a realistic range spanning low to one year beyond the replacement age.

The empirical values used to initialize these household attributes are primarily based on large-scale surveys conducted by ANWB and Paradies, offering valuable insight into Dutch travel behaviour and mobility preferences. While these sources provide robust average estimates, limitations arise due to potential biases such as those introduced by the intention-behaviour gap and the lack of reported variance measures. Consequently, standard deviations for input distributions had to be assumed.To account for this uncertainty and explore the impact of these assumptions, the corresponding parameters were well incorporated into the sensitivity analysis.

Household Differentiation
To capture not only small within-neighbourhood heterogeneity but also larger disparities between neighbourhoods, which can be seen as socio-spatial demographic groups, archetypes are identified.

The number of archetypes was determined using the elbow method in a k-means clustering process (Figure 4.4a), selecting the point where additional clusters yielded diminishing returns in within-cluster variance reduction. Input features: average income and the percentage of WO or HBO-educated inhabitants, were standardized and drawn from CBS

Table 4.2.: Overview of Household Initialization Parameters

Parameter	Value (default → range)	Description
Ownership and Preferences		
`ownership`	Probabilistic → {CV, EV, NC}	Initial ownership status, based on `perc_EV` and `perc_NC`
`no_car_preference`	0.81 → {0, 1}	81% of NC agents have intrinsic no-car preference
Mobility and Budget Needs		
`daily_km`	$\sim \mathcal{N}(23, 46) \to [5, \infty)$	Actual daily driving demand in km [ANWB, 2025]
`range_need`	$\sim \mathcal{N}(461, 0.5 \cdot 461)$	EV range requirement [ANWB, 2025]
`car_budget_EV_h`	$\sim \mathcal{N}(29,036, 7259)$	EV new purchase budget [ANWB, 2025]
`car_budget_EV_l`	$\sim \mathcal{N}(12,997, 3249)$	EV occasional budget [ANWB, 2025]
`car_budget_CV_h`	$\sim \mathcal{N}(24,894, 6223)$	CV high-end purchase budget[ANWB, 2025]
`car_budget_CV_l`	$\sim \mathcal{N}(10,749, 2687)$	CV low-end purchase budget [ANWB, 2025]
`max_running_cost`	$\sim \mathcal{N}(200, 50)$	Max affordable monthly vehicle cost
Behavioural Characteristics		
`ev_knowledge`	0.70 → [0, 1]	Initial perceived sufficient ability/awareness of EVs [Paradies et al., 2023]
`routine_buyer`	0.34 → [0, 1]	Fixed-type preference bias [Paradies et al., 2023]
Vehicle Lifespan		
`car_replacement_age`	$\sim$ lognormal($\mu = \log(6.6), \sigma = 0.6$)	Expected lifespan before replacement [Paradies et al., 2023]
`current_car_age`	Uniform(0, `car_replacement_age` +1)	Assumed initial car age
Demographic characteristics		
`diff_inc`	1.0 → [0, 2]	Income deviation scalar from national average
`diff_dens`	1.0 → [0, 2]	Housing Density deviation scalar from national average
`diff_educ`	1.0 → [0, 2]	Education (HBO or WO) level deviation scalar
`arch`	0,1,2,3,5	Archetype number, descripives can be found in 4.2.2, 5=no specific arch

district-level data ([Municipality The Hague, 2024]). The living density situation was based on research by TNO, which identified different infra-living neighbourhoods [Kaas et al., 2024]. Each resulting cluster can be interpreted as an archetype and characterized using the mean values of the input variables, visualized in a radar plot (Figure 6.9a).

These archetypes are spatially mapped onto The Hague's districts (Figure 6.9b) and linearly linked to agent attributes through neighbourhood-level scalars: income (`diff_inc`), education (`diff_educ`), and housing density (`diff_dens`). These scalars modify key agent decision processes such as EV affordability evaluations, knowledge levels in the CODEC framework, and access to infrastructure.

Within each archetype, agents retain stochastic variation to reflect intra-neighbourhood diversity. This layered structure captures both structural inequality and behavioural nuance, forming the foundation for the model's policy and scenario analysis.

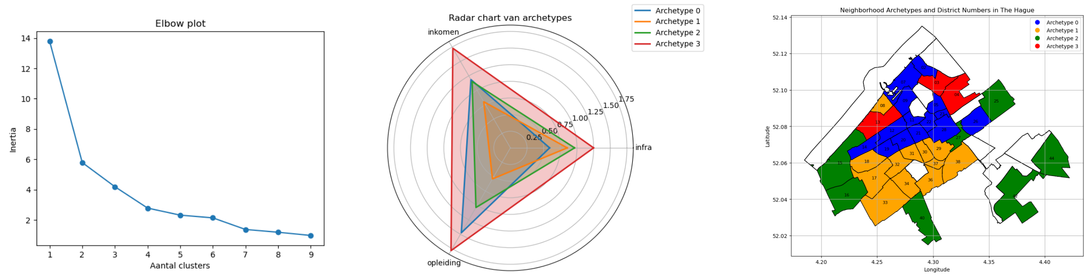

(a) (b) (c)

Figure 4.4.: (a) Elbow plot for cluster selection. (b) Radar plot of archetype features. (c) Spatial distribution across The Hague.

Household Advancement

At each simulation step, agents update their internal state through the `step()` function, following a structured decision process. First, the agent increments its current vehicle age by one year. It then checks whether it has entered the decision cycle via the attention phase: this occurs when the vehicle age exceeds the pre-set replacement threshold. Routine buyers of combustion vehicles (CVs) bypass detailed evaluation and automatically renew their vehicle, resetting the age without entering the decision process.

If the agent is in the decision phase and currently does not own a car (NC), it may either maintain this status, which it does if it has a strong intrinsic preference (`no_car_preference` = 1) or reconsider ownership with a 50% probability if this preference is weaker, reflecting the option of those into forced carless state to get out of this.

For agents who do own a vehicle, research by Mashhoodi and Van Der Blij [2020] on Amsterdam indicates that 16% of residents in urban areas seriously consider going car-free. It is assumed that half of these actually follow through, and that this share increases by 2% annually to align with broader no-car mobility trends. This is modelled as:

$$P_{\text{switch to NC}} = 0.08 \times (1.02)^t \tag{4.1}$$

where t is the current simulation year.

If the agent considers purchasing a vehicle, it evaluates four vehicle types: high- and low-tier electric vehicles (`EV_h`, `EV_l`) and combustion vehicles (`CV_h`, `CV_l`). Each option is scored by multiplying the enablement and intention phases, both derived from the CODEC decision framework:

$$\text{score}_v = \text{Enablement}_v \times \text{Intention}_v,$$

where v indexes the vehicle type.

The enablement phase measures whether the agent *can* buy the vehicle, combining several probabilistic components. These include financial feasibility, computed as the cumulative distribution function (CDF) of the budget relative to vehicle price:

$$\text{invest\_score} = \Phi\left(\frac{\text{budget} - \text{price}_v}{0.5 \times \text{price}_v}\right),$$

technological fit based on range needs relative to vehicle range:

$$\text{range\_score} = 1 - \Phi\left(\frac{\text{range\_need} - \text{range}_v}{0.5 \times \text{range}_v}\right),$$

and, for EVs specifically, factors like EV knowledge level (ev_knowledge), infrastructure availability ($ev\_infra\_coverage$), and second-hand market availability (adjusted by $ev2avail$ for low-tier EVs). The enablement score is the product of these factors:

$$\text{Enablement}_v = \text{range\_score} \times \text{invest\_score} \times \text{knowledge\_score} \times \text{infra\_score} \times \text{availability\_score}.$$

Note that all scores lie in the interval $[0, 1]$, where higher values indicate more favourable conditions. The investment score reflects the probability that the vehicle is affordable (i.e. the household's budget exceeds its price), whereas range score captures the probability that the vehicle's range exceeds the household's range need. A low value in any single component proportionally reduces the overall enablement score, mirroring the conjunctive nature of the decision process.

The intention phase reflects how much the agent *wants* the vehicle. It combines the attractiveness of the purchase investment and running costs through their respective CDFs:

$$\text{atrr\_invest\_score} = \Phi\left(\frac{\text{budget} - \text{price}_v}{0.5 \times \text{price}_v}\right),$$

$$\text{atrr\_run\_score} = \Phi\left(\frac{\text{max\_running\_cost} - \text{running\_cost}_v}{0.5 \times \text{running\_cost}_v}\right).$$

Non-monetary vehicle attributes such as sustainability, refuelling convenience, driving pleasure, and innovativeness are weighted based on empirical ratings from [Paradies et al., 2023] and adjusted for vehicle tier according to 2023 TNO research on the second-hand EV market.

Social influence plays a key role: conformity increases preference for vehicle types common among the agent's neighbours, quantified by the normal CDF of the local share of neighbours owning the same vehicle type. Distinction drives some EV owners to prefer standing out when EV adoption is low, modelled as one minus the normal CDF of the share of neighbours owning the same EV type.

The final intention score sums these components with weights (again from the [Paradies et al., 2023] survey) reflecting their relative importance:

$$\begin{aligned}\text{Intention}_v = \; & 0.14 \times \text{atrr\_invest\_score} + 0.10 \times \text{atrr\_run\_score} + 0.42 \times \text{feature\_score} \\ & + 0.08 \times \text{distinction\_score} + 0.25 \times \text{conformity\_score}.\end{aligned}$$

After calculating scores for all vehicle types, the agent selects a vehicle probabilistically weighted by these scores. If no vehicle meets a positive threshold, the agent either keeps its current car or opts for no car, based on annual travel distance and chance.

This decision process embeds social norms dynamically by considering neighbouring households' ownership types, creating a feedback loop between local social structure and vehicle technology diffusion.

4.2.3. Agent: Government Policy Mechanisms

The agent-based model incorporates an institutional agent, the Government, that can implement policy interventions to influence electric vehicle (EV) adoption dynamics. This component is encapsulated in the `Government` class and is responsible for applying and reversing a set of predefined policy levers at specified simulation time steps.

This flexible policy mechanism allows for dynamic experimentation with different government interventions over time. It supports evaluation of both individual and combined effects of demand-side (e.g., subsidies, marketing) and supply-side (e.g., infrastructure, regulation) policies on EV adoption trajectories.

The government module functions as a policy scheduler and executor. Its key attributes and methods are as follows:

- **Policy Configuration:** At each time step, the model checks if policies are scheduled via the `timed_policy_configs` dictionary. If active policies exist, the government applies them via the `apply_policies()` method.
- **Policy Reversibility:** To ensure policy effects are non-cumulative unless desired, previously applied policies are undone before applying new ones using the `undo_policies()` method.

At each designated time step t, the following sequence is executed:

1. The model checks if $t \in$ `timed_policy_configs`.

2. The government first invokes `undo_policies()` to revert the effects of the previous round of policies.
3. It then applies the newly defined set of policies using `apply_policies()`.
4. The policies active at t are stored in `last_applied_policies` for reference and reversibility.

This remainder of this section outlines the policy levers implemented in the agent-based model, designed in alignment with the theoretical findings discussed in Chapter 2. The model simulates four core policy interventions—*targeted subsidies*, *infrastructure expansion*, *community information campaigns*, and *low-emission regulation*. Each targets key dimensions of EV adoption: affordability, range confidence, behavioural awareness, and normative shift.

1. Targeted Subsidy

To address affordability barriers and accelerate adoption among lower-budget households, a temporary €2000 subsidy is granted to agents whose EV purchase budget falls short of the market price for a second-hand EV plus a €1000 margin. This heuristic proxy identifies cost-sensitive buyers who are likely to face financial constraints in affording an EV.

Formally, for each agent a, the applied subsidy $S(a)$ is defined as:

$$S(a) = \begin{cases} 2000, & \text{if budget}_{EV_l}(a) < P_{EV_l} + 1000 \\ 0, & \text{otherwise} \end{cases}$$

Here, $\text{budget}_{EV_l}(a)$ is the agent's disposable car budget for a second-hand EV, and P_{EV_l} is the market price of that vehicle type. The effective price is adjusted at the point of decision-making by modifying the agent's `car_subsidy` attribute, which then ups their budget throughout the model. When the subsidy policy is removed `car_subsidy` becomes 0 again.

While many real-world subsidy schemes offer around €1000 per household, this model assumes a higher amount of €2000. The rationale is that this more generous subsidy could realistically help bridge the affordability gap for lower-income adopters. Since the policy is tightly targeted and temporary, applying only to those with insufficient EV budgets and only during a designated policy step, the total financial burden remains limited whilst the policy impact on equity and adoption potential is maximized. This approach aligns with the literature on equity-focused EV policy design [Hopkins et al., 2023; Joshi et al., 2022], while also incorporating a budget-responsive targeting mechanism.

2. Infrastructure and Access Expansion

To simulate improvements in public charging infrastructure and corresponding behavioural effects, infrastructure coverage is increased heterogeneously across space. Underserved zones (low charger density or low uptake) receive a 30% boost in public access, which is should an attainable (under grid constraints) and sufficient (under low anxiety) goal according to [Mashhoodi and Van Der Blij, 2020] , while well-equipped areas receive a 15% improvement. Additionally, with the total amount of infra going up to sufficiency the range anxiety among residents of all zones is reduced by 20%, consistent with behavioural responses to charging security in the literature [Wang et al., 2023].

Formally, for an agent a, the perceived infrastructure coverage is updated as:

$$C^*(a) = C + \delta_{\text{infra}}(a)$$

$$\delta_{\text{infra}}(a) = \begin{cases} 0.30, & \text{if } \texttt{diff\_dens}(a) < 0.9 \quad \text{(underserved area)} \\ 0.15, & \text{otherwise} \end{cases}$$

The agent's adjusted perceived EV range requirement becomes:

$$R^*(a) = R(a) \times (1 - \delta_{\text{range}}) \quad \text{with} \quad \delta_{\text{range}} = 0.15$$

These adjusted values affect both the vehicle eligibility filter and the utility score in the decision-making process. Note that the infrastructure expansion effect is not removed once implemented, as it reflects structural improvements to the built environment rather than temporary incentives.

3. Community Information Campaigns

Based on findings that knowledge gaps and peer influence are crucial in shaping EV perceptions, the model implements an outreach campaign targeting hesitant or disadvantaged agents based on their neighbourhood [Song and Potoglou, 2020; Varghese et al., 2024].

Agents with low income and low educational attainment receive a 60% increase in EV-related knowledge, while others receive a 30% increase due to passive exposure. Additionally, perceived range needs are reduced by 25% for all agents to reflect growing awareness of typical daily driving patterns and improved infrastructure information.

Let $K(a)$ represent the agent's knowledge and $R(a)$ their perceived range need. These are adjusted as:

$$K^*(a) = K(a) \times (1 + \delta_K), \quad R^*(a) = R(a) \times (1 - \delta_R)$$

where $\delta_K = 0.60$ for agents residing in lower-income and lower-education neighbourhoods (operationalized for archetype 1 and 2), and $\delta_K = 0.30$ for all other agents. The parameter $\delta_R = 0.25$ is applied uniformly across the population. Information campaigns can be implemented multiple times, both formally and realistically, as their effects are assumed to be cumulative; repeated exposure enhances learning and retention. Furthermore, it is assumed that the gains in knowledge and reductions in range anxiety are persistent within the simulation time frame and do not decay over time.

4. Zero-Emission Zone Regulation

To model regulatory mechanisms, the policy simulates the introduction of a zero-emission zone (ZEZ) that restricts the use of internal combustion engine (ICE) vehicles. Two implementation modes are considered: a mandate-enforced version for neighbourhoods within the city centre boundaries, and a non-enforced (soft) version for other neighbourhoods, where the effects are experienced only indirectly

In the mandate-enforced version, all conventional vehicle (CV) owners in the ZEZ who are likely to drive older emission class vehicles (estimated at 50%, as per Central Bureau Statistics Netherlands CBS [2023]) are forced to reconsider their vehicle ownership. These agents are required to re-evaluate their mobility needs immediately by setting their vehicle age to the replacement threshold and disabling routine decision-making. Additionally, ICE attractiveness is decreased to reflect stricter norms. Formally:

$$\text{If } \texttt{ownership}(a) = \text{CV and rand}() < 0.5: \quad \begin{cases} \mathit{routine\_buyer}(a) = \texttt{False} \\ \texttt{current\_car\_age}(a) = \texttt{car\_replacement\_age}(a) \\ \texttt{cv\_utility\_discount}(a) = 0.6 \end{cases}$$

In the non-enforced version, no mandates are applied, but behavioural shifts are encouraged through altered cost-benefit perceptions. EVs become more publicly salient and ICE vehicles less attractive due to anticipated access limitations in city centres. Thus, all agents—regardless of location—perceive a reduced utility from ICE vehicles, but a it less: `cv_utility_discount`(a) is set to 0.8. This results in both direct enforcement and soft signalling mechanisms contributing to accelerated behavioural shifts, consistent with empirical findings from Gao et al. [2022]; Lee et al. [2021]. The zero-emission zone regulation, for a realistic portrayal, is assumed to persist when implemented at an earlier timestep. Its effects are applied once and not reversed over time, as they represent a behavioural change rather than a temporary policy.

5. Model Validation and Exploration

5.1. Verification and Validation

To ensure that the model produces trustworthy results, various verification techniques were applied to confirm the correct implementation of mechanisms, the absence of bugs, and appropriate boundary behaviour. Following verification, further validation using real-world data and literature was conducted to assess whether the model's outputs are meaningful and whether the model is indeed *fit for purpose*. Given the model set-up presented in Chapter 4, validation is performed by layer: starting with the agent level, followed by the system level, and finally the policy implementation.

5.1.1. Agent-Level

Agent Verification: Decision Logic and Edge Cases

During early development, individual household agents were equipped with logging tools to monitor internal state transitions (e.g., vehicle ownership), decision triggers (e.g., available subsidies, peer influence), and final choices. This facilitated close inspection and debugging of agent behaviour under diverse conditions.

Controlled test cases of extreme conditions were used to verify agent decision logic. Boundary cases confirmed that under what should be perfect conditions, adoption was effectively inevitable, while under highly constrained conditions, agents consistently refrained from adoption. Edge cases, such as agents flagged as routine buyers, were also tested and correctly bypassed the CODEC decision-making process, maintaining their existing vehicle type regardless of context. These checks confirmed the correct implementation of both standard and exceptional behavioural rules.

Agent Validation: Micro-Level Behavioural Realism

Individual agent behaviour was qualitatively compared with real-world behavioural tendencies reported in the literature. For instance, high-income agents in the model tended to adopt high-end EVs—a pattern supported by empirical research ([Axsen and Kurani, 2013]). Similarly, agents with long daily travel distances often rejected EVs due to range limitations, reflecting real-world concerns around range anxiety and infrastructure gaps ([Gnann et al., 2023]). These findings confirm that agent decisions not only function correctly but also align with plausible behavioural expectations.

5.1.2. System-Level

System Verification: Interaction Dynamics and Execution Logic

A simplified 25-agent "play model" on a 5×5 grid was used to test multi-agent interactions, particularly peer influence dynamics. Conformist agents reliably adopted EVs when surrounded by adopters, and no unrealistic behaviours (e.g., cyclic adoption or instability) emerged. Larger-scale tests (up to 2,000 agents) under varying demographics, grid sizes, and random seeds confirmed correct sequencing of model events: policy activation, household decisions, and vehicle updates.

In the base case as shown in Figure 5.1, the percentage of EV-owning households eventually surpasses that of CVs, but due to limited EV uptake among routine buyers, adoption stagnates and later declines as the non-car-owning (NC) trend grows. While this outcome is debatable in terms of realism as this is a big uncertainty, it only occurs beyond the planned simulation timeline, and future routine buyer behaviour is uncertain. As such, it does not affect the model's validity for its intended purpose. Lastly minor variations in spatial clustering were observed but had no meaningful effect on outcomes, confirming execution stability and logical consistency.

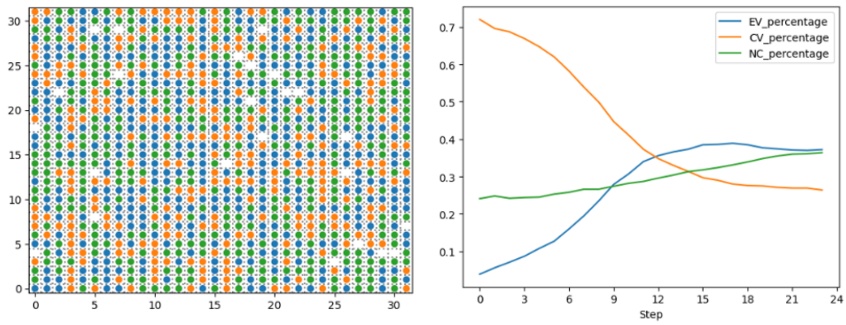

Figure 5.1.: Playmodel Shot: Verification of Spatial Grid and Emergent Behaviour

System Validation: Macro-Level Trends and Historical Comparison

The model was validated at the macro level by running simulations with 1,000 agents over 1,000 iterations, extending an extra ten years beyond the planned analysis timespan. The results produced an S-shaped adoption curve, typical of diffusion-of-innovation processes (see Figure 5.2a) ([Rogers, 1983]; [Sierzchula et al., 2014]). Historical EV fleet data from 2022–2025 ([Central Bureau Statistics Netherlands CBS, 2023]) was used to compare adoption trajectories, which fell within the model's standard deviation range (see Figure 5.2). While not conclusive, this alignment is a positive indication of the model's external validity.

To further validate the model structure and heterogeneity implementation, a series of runs was conducted across the four defined neighbourhood archetypes. As shown in Figure 5.2b, adoption trajectories diverged in plausible and expected ways: higher-income, better-educated, and well-infrastructured archetypes adopted more rapidly than those with lower income, education, and infrastructure levels. This confirms that the implemented differences between archetypes meaningfully affect outcomes and align with known socio-demographic adoption patterns reported by Central Bureau Statistics Netherlands CBS [2023]. These results also align with Rogers [1983], as one can observe different segments of the S-curve emerging across archetypes, suggesting that they represent distinct adopter categories within the diffusion of innovation framework.

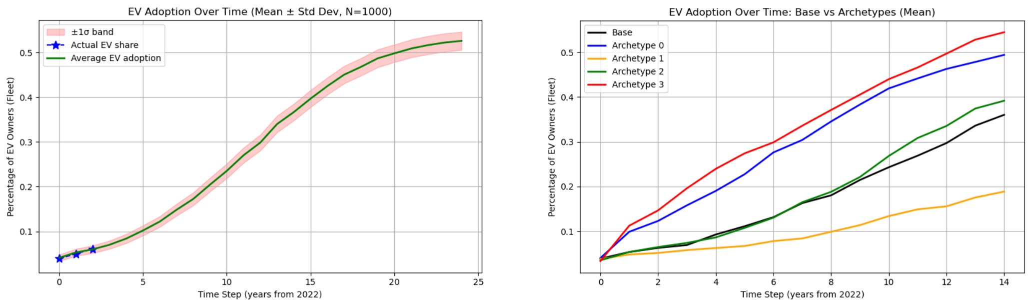

(a) (b)

Figure 5.2.: (a) Simulated vs. Historical EV Adoption (2022–2025). (b) Simulated Baseline vs. Archetype Adoption Patterns (2022–2035)

System Robustness: Sensitivity Analysis

To assess robustness under uncertainty, a sensitivity analysis was performed by perturbing key input parameters by ±10%, both individually and all at once. As shown in Figure 5.3, the model is most sensitive to EV price and range, which is expected given their central role in adoption decisions. Feature weights also strongly influence outcomes, as they shape how agents compare EVs and CVs. This is amplified for the driving pleasure and sustainability, which score very differently across vehicle types.

Second-hand EV availability factor also plays into such a difference as particularly occasional buyers depend on affordable used options. Routine buyers, who are almost resistant to system change, make sense to slow the transition significantly. Most other inputs showed limited sensitivity, suggesting that assumptions made during model formalisation (e.g., standard deviations) do not distort results. Overall, the model responds predictably to input variation, with outcomes remaining stable and within plausible bounds.

5.1.3. Policy-Level

Policy Verification: Trigger and Implementation Logic

To verify correct policy behaviour, each mechanism (e.g., subsidies, zero-emission zones, information campaigns) was tested in isolation using both simulation and logging. Logging confirmed, for example, that infrastructure expansion dynamically updated agent enable scores. As shown by the simulated timelines in Figure 5.4, all policy scenarios begin to diverge precisely at t=4, the intended activation point. Before this, runs follow identical trajectories, confirming correct seed control, accurate trigger timing, and persistent effects over time.

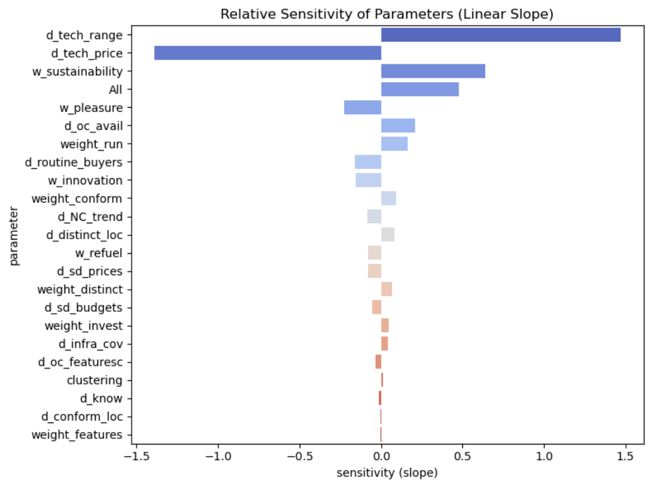

Figure 5.3.: Parameter-Wise and Combined Perturbation Effects (±10%) on EV Adoption Outcomes

Delayed adoption responses reflect real-world lags due to consumer decision timelines, vehicle availability, and peer influence [Li et al., 2023]. Interestingly, policies initially increase divergence from the baseline but later converge again, suggesting they partly accelerate adoption rather than alter long-term outcomes, which is a realistic result.

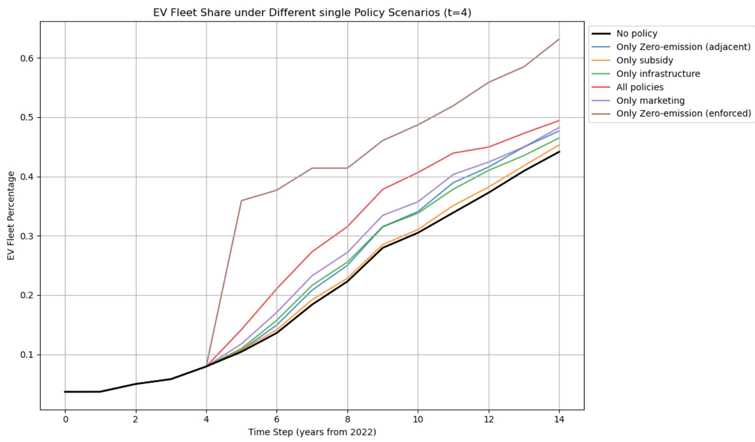

Figure 5.4.: Policy implementation

Policy Validation: Realism and Interpretability of Policy Effects

Despite their simplified implementation, the modelled policy scenarios produce adoption patterns that align well with empirical expectations and behavioural theory, indicating that they are fit for purpose. For instance, income-targeted subsidies primarily benefit lower-income households which is consistent with findings that affordability is a key barrier for this group, but looking overall they have limited effect [Joshi et al., 2022; Siebenhofer et al., 2021]. The sharp increase in adoption following the introduction of enforced zero-emission zones mirrors the kind of disruption needed to break habitual repurchasing behaviour [Gao et al., 2022]. Likewise, the early surge in the marketing-only scenario reflects the power knowledge that is came up time and time again in the literature [Sikder et al., 2023].

These dynamics suggest that the model appropriately captures the mechanisms through which different types of interventions influence EV uptake, particularly across socio-demographic groups. This also being supported by face validation from TNO mobility experts, it can be reasonably concluded that it provides a suitable basis for exploring how policies can affect both the speed and distributional fairness of the transition to electric mobility.

5.2. Exploration

5.2.1. Scenario and Uncertainty Exploration

To identify scenarios with high output variance in my experimental setup, a Sobol global sensitivity analysis was performed. Given the complexity of the agent-based model (ABM) and the presence of many interacting and potentially correlated uncertainties, the Sobol method is well-suited. Unlike simpler feature scoring methods or scenario-focused tools like PRIM, Sobol provides a comprehensive, variance-based global sensitivity analysis. It quantifies both the individual (first-order) and interaction (total-order) effects of each input on model output variance. This is crucial for large ABMs, where interactions between parameters can have substantial, non-linear impacts. By capturing these effects, the Sobol method enables the identification of input combinations that most influence output variability, making it an ideal foundation for selecting particularly informative or extreme scenarios for further investigation.

The analysis focused on ten behavioural and technological parameters expected to affect electric vehicle (EV) adoption most given the small sensitivity analyses of Section 5.1.3. The variables are now perturbed more, with ranges going up 50% up and down. Convergence analysis (see Figure 5.5a) shows that Sobol indices stabilize from around 4,000 samples onward, indicating robustness in the sensitivity estimates. As such, the results at 5,000 samples are used as the basis for further interpretation and can be found in Figure 5.5b.

The first thing to notice is that this presents quite a different picture compared to the validation sensitivity analyses. This can be explained by the broader scope of variation and the presence of non-linear interactions, which imply that significant effects may only occur once certain thresholds are crossed. Looking at the new results, the total-order sensitivity indices (Figure 5.3) reveal that the most influential parameters are:

- `w_pleasure` – the weight assigned to driving enjoyment in household decision-making,
- `d_oc_avail` – the average availability of local charging infrastructure.

Both parameters exhibit high total-order effects and relatively large gaps between their total-order and first-order indices. This implies that their influence is not only direct, but also significantly mediated by interaction with other parameters. For example, driving enjoyment may reduce EV uptake more strongly when charging availability is limited, or when

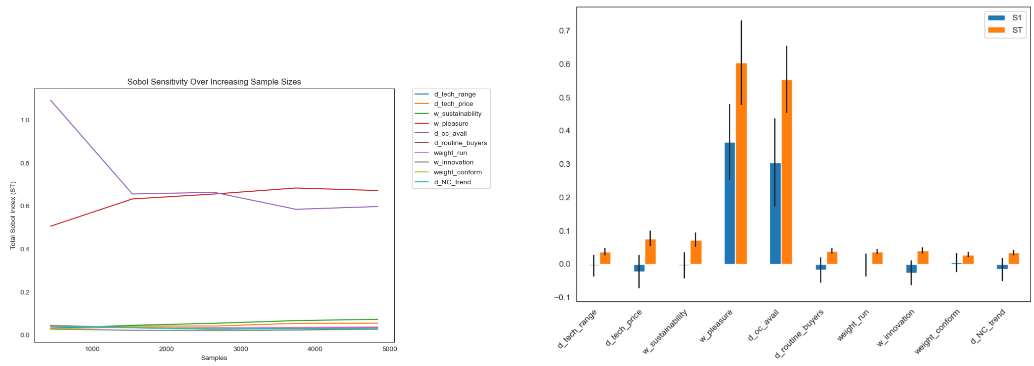

(a) (b)

Figure 5.5.: (a) SOBOL global sensitivity analyses convergence. (b) SOBOL global sensitivity analyses parameter outcomes.

vehicle prices remain high. This highlights also the importance of non-linear and conditional dynamics in the model and the need to test further experiments also with combined scenarios.

The rest of the parameters show to have meaningful total-order sensitivity, though their first-order effects are often smaller or more uncertain, with several confidence intervals including low. This suggests that these parameters likely contribute to outcome variation through interactions or non-linear effects, rather than direct, isolated influence.

To guide the structured uncertainty exploration in the EMA Workbench, six parameters were selected based on a combination of model sensitivity results and real-world interpretability. The selection followed a three-step logic. First, we included the two most influential parameters from the Sobol global sensitivity analysis: `w_pleasure` (the weight assigned to driving enjoyment) and `d_oc_avail` (the average availability of local charging infrastructure). These variables showed the highest total-order effects and are thus central to output variance. Second, we added the two most prominent parameters from the earlier small sensitivity analyses: `d_tech_price` (EV price development) and `d_tech_range` (EV range development). These are not only impactful according to Sobol, but are also widely emphasized in the literature as key drivers of EV adoption. Third, from the remaining Sobol parameters, we selected two additional variables that, while slightly less dominant in sensitivity magnitude, still have notable influence and represent different real-world mechanisms. `d_NC_trend` reflects the background trend in conventional vehicle attractiveness, and `d_routine_buyers` captures behavioural lock-in through routine-based decision-making.

Together, these six parameters (`w_pleasure`, `d_oc_avail`, `d_tech_price`, `d_tech_range`, `d_NC_trend`, and `d_routine_buyers`) form a balanced and diverse set for scenario exploration. Their definitions and value ranges are listed in Table 5.1. This approach ensures that the scenario space is shaped by uncertainties that are both impactful in the model and relevant to policy and behavioural dynamics.

5.2.2. Policy Exploration

To also identify a broad and behaviourally diverse set of high-performing policy configurations for further testing under uncertainty, a structured policy discovery approach was implemented across the four predefined socio-spatial archetypes and the base archetype which is the average of the Netherlands. This involved generating a wide space of temporally structured policy schedules by combining four intervention types, subsidies, infrastructure investment, marketing, and zero-emission mandates, across three decision points. Logical constraints were applied to ensure consistent policy progression, most notably by enforcing the persistence of the zero-emission mandate once introduced.

Given the combinatorial nature of the policy space, an exhaustive evaluation of all possible schedules was computationally impractical. Therefore, a strategically sampled subset of 213 scenarios was selected to balance representativeness and tractability. This included: (i) the base scenario without any policies; (ii) the full set of single-policy schedules, where each of the four policies (subsidy, infrastructure, marketing, zero-emission) was implemented individually at one of the three decision points ($3 \times 4 = 12$ scenarios); (iii) a stratified sample of 50 schedules for each of four distinct timing patterns of the zero-emission policy namely: never introduced (AAA), introduced at $t = 8$ (AAB), at $t = 6$ (ABB), or from $t = 4$ onward (BBB). From each of these zero-emission patterns a random variation in other co-occurring policies is sampled, but atleast 20 is set to have atmost 3 policies implemented over the years, to make sure enough samples are kept which are relatively easily implementable.

Figure 5.6 shows how adoption trajectories vary across archetypes. Many of the policy combinations lead to improvements for most archetypes, with the better-performing combinations increasing EV fleet share by approximately 10 percentage points. Notable by the y-shape, for Archetypes 1 and 3, some policies are highly effective while others have little to no impact. This aligns with the fact that these archetypes are relatively wealthier, and therefore benefit less from policies specifically targeted at lower-income or disadvantaged groups.

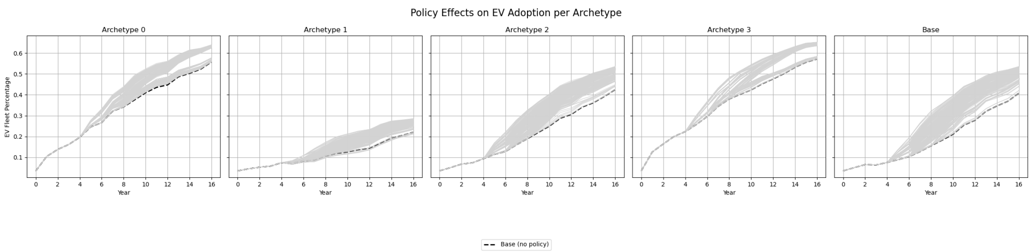

Figure 5.6.: EV Adoption under Policy Scenarios across Archetypes. Grey lines: individual policies; dashed line: base scenario.

Performance–equity trade-offs were evaluated using the mean final EV share and the standard deviation across archetypes, applying equal (50/50) weights for simplicity. The resulting scatter is shown in Figure 5.7. In this figure, an ideal policy would appear in the top-left corner, combining high overall adoption with low inequality across archetypes. However, even with relatively ambitious literature-based policies, few policy-combinations reach this region, highlighting a clear trade-off between equity (low standard deviation) and performance (high final EV share). As expected, less constrained policy combinations tend to perform better overall. Nonetheless, the figure also shows that certain single policies or those with three or fewer policies over the years, can already lead to significantly improved outcomes on either or both metrics.

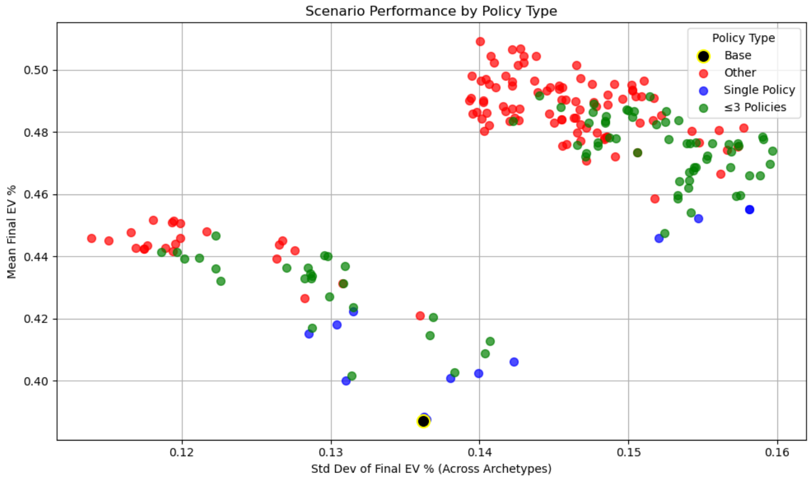

Figure 5.7.: Left: Adoption vs. inter-archetype variance. Right: Clustering of top-performing policy sequences.

Using the K-means algorithm, PCA clustering of the 25 top-performing scenarios (Appendix B) revealed distinct timing patterns across different planned policy combinations. Figure 5.8 presents clustered heatmaps of policy activation patterns across scenario groups. Several key insights emerge. Zero-emission zones are extremely dominant when allowed, but they are generally most effective when implemented at time step $t = 6$ or later. This is likely because the population is then "more ready" to actually choose an EV. In many patterns, this readiness is supported primarily by earlier deployment of marketing, infrastructure, and subsidies. Notably, when only three policies are allowed, the clusters reveal that marketing and infrastructure are the most consistently effective options to enable this transition.

When zero-emission zones are excluded—an interesting scenario given their controversial and forceful nature—other policy combinations can still yield strong performance. In these cases, multiple policies are often activated simultaneously at the same time step, likely to provide a strong cumulative boost. Marketing and infrastructure tend to be more effective when introduced earlier, with infrastructure ideally maintained over a longer period. Subsidies appear most effective when timed at both the beginning (to encourage early adopters) and the end (to bring in those who are nearly ready).

These findings highlight that well-targeted and well-timed policy mixes can lead to high EV adoption outcomes, even under limited intervention conditions.

5.2.3. Selected Policy Configurations for Further Exploration

Based on the analysis of performance–equity trade-offs and policy timing clusters, six policy configurations were selected for structured scenario exploration. These configurations reflect diverse strategic logics, ranging from minimalist and phased approaches to high-intensity packages, and allow comparison of different timing, targeting, and complexity

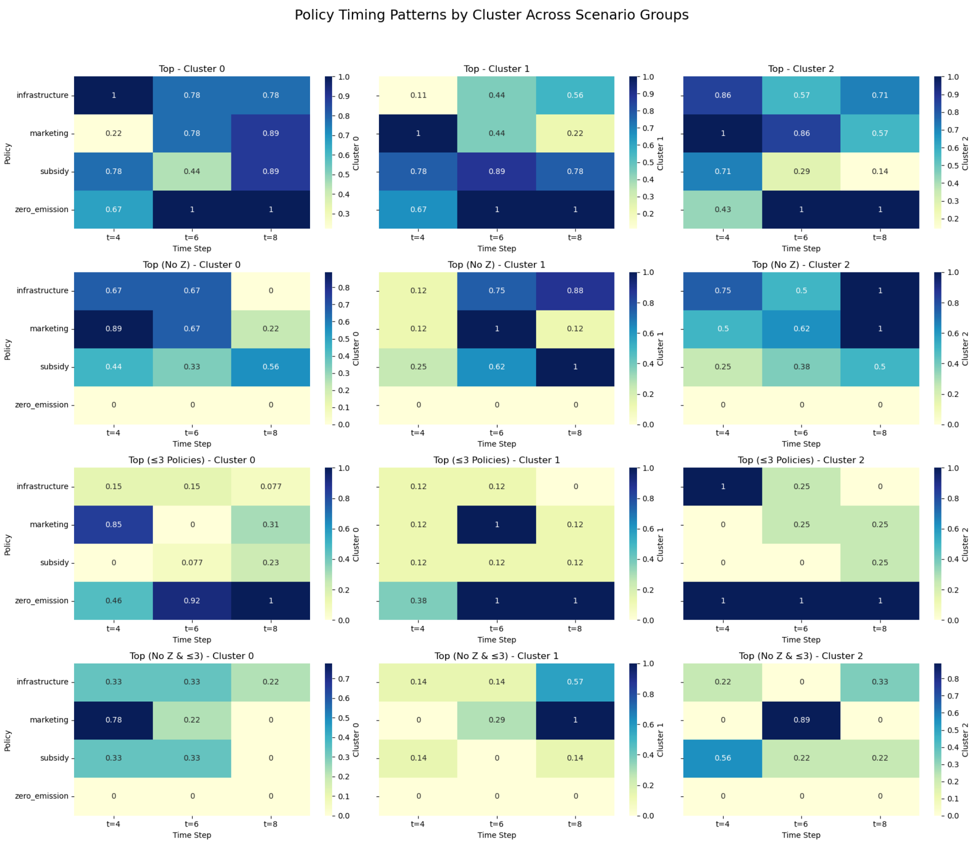

Figure 5.8.: Policy timing frequency among top- and bottom-performing clusters.

levels. Each Policy Combination (PC) includes specific activation timings for the four key policy levers: infrastructure investments (I), marketing campaigns (M), purchase subsidies (S), and zero-emission zones (L). The selection is intended to span both high-performing cluster archetypes and theoretically relevant alternatives to support robust policy evaluation.

- **PC1. Baseline (No Policy):** This scenario serves as a reference case and includes no interventions.
  Policy Configuration: {4: None; 6: None; 8: None}
- **PC2. Infrastructure + Marketing, Sustained:** A light-touch but consistent mix focusing on accessibility and awareness, without regulation or subsidies.
  Policy Configuration: {4: I, M; 6: I, M; 8: I}
- **PC3. Early Subsidy with Infrastructure:** A cost-relief–driven strategy where financial incentives are paired with infrastructure, fading out toward the end.
  Policy Configuration: {4: I, S; 6: S; 8: I}
- **PC4. Delayed ZEZ with Subsidies and Awareness:** A politically realistic phasing that begins with soft levers and ends with regulation, leaving infrastructure to the market.
  Policy Configuration: {4: M, S; 6: I, Z; 8: Z}
- **PC5. Infrastructure and ZEZ:** A focused mix of physical and regulatory levers, where infrastructure and marketing pave the way for regulation.
  Policy Configuration: {4: I; 6: I, M; 8: Z}
- **PC6. Heavy Package – All Policies:** A maximal push with every policy lever activated early and maintained throughout.
  Policy Configuration: {4: I, M, S, Z; 6: I, M, S, Z; 8: I, M, S, Z}

5.2.4. Experimental Design

This study uses agent-based simulation to evaluate the performance of electric vehicle (EV) policy interventions under behavioural and technological uncertainty. The experimental setup addresses three key research objectives: (i) measuring adoption speed and total uptake, (ii) assessing equity and inclusion across socio-economic groups, and (iii) identifying robust strategies under plausible future developments.

Design Set-up

The experimental design manipulates two main categories of independent variables:

- *Policy configuration:* Six policy bundles are tested, including both zero-emission zone (ZEZ) and non-ZEZ alternatives. These combinations were selected based on prior scenario exploration and clustering of top configurations (see Section 5.2.2). They vary in policy type (e.g., marketing, subsidy, infrastructure, ZEZ) and timing (activation in 2026, 2028, or 2030).

- *Uncertainty parameters (continuous ranges):* Instead of fixed future scenarios, uncertainty is introduced through continuous variation in key behavioural and technological parameters, identified via Sobol sensitivity analysis (see Section 5.2.1). These parameters define the ”scenario space” across which policies are evaluated using the EMA Workbench.

This enables robustness testing of policy performance across diverse and realistic system conditions. Policy–uncertainty combinations are tested in two progressively realistic urban environments: first, simplified runs across archetype neighbourhoods (as constructed in Section 4.2.2) are used for broad and deep analyses of policy combination performance, allowing for the selection of optimal configurations; second, a realistic spatial layout of 44 districts in The Hague, incorporating empirically informed heterogeneity, enables spatially explicit policy insights for the actual implementation of the final policy package(s). This setup balances experimental control with real-world policy relevance. This setup balances experimental control with real-world policy relevance.

Scenario–Policy Sampling

Using the EMA Workbench, a full-factorial policy design is embedded within a sampling framework that varies the uncertainty parameters described above (see Table 5.1). Each of the 6 policy configurations is evaluated across 1,000 Latin Hypercube samples from the uncertainty space, resulting in thousands of unique policy–context combinations. This approach ensures broad and efficient coverage of plausible futures while preserving statistical diversity in the experimental design, which is why LHS is chosen over the more random methods such as Monte Carlo sampling.

This enables the evaluation of how each policy performs under a wide range of assumptions, the identification of conditions under which a policy succeeds or fails, and the screening of interventions that are robust and broadly effective.

Table 5.1.: Uncertainty–Policy Matrix: Parameters Sampled Within Ranges Across All Final Policy Configurations

Uncertainty Parameter (Range)	PC1 {4: None, 6: None, 8: None}	PC2 {4:I,M; 6:I,M; 8:I}	PC3 {4:S,I; 6:S; 8:I}	PC4 {4:M,S; 6:I,Z; 8:Z}	PC5 {4:I; 6:I,M; 8:Z}	PC6 {4:I,M,S,Z; 6:I,M,S,Z; 8:I,M,S,Z;	**Interpretation (sampled range)**
`w_pleasure` [0.1–1.5]	X	X	X	X	X	X	Preference for driving fun; high = emotional EV resistance
`d_NC_trend` [0.1–2]	X	X	X	X	X	X	Reverting to no car lifestyle; high =shared mobility uptake
`d_routine_buyers` [0.1–0.6]	X	X	X	X	X	X	Share of habit-driven buyers; high = slow change
`d_oc_avail` [0.1–1.0]	X	X	X	X	X	X	Access to charging; low = infra bottlenecks
`d_tech_price` [0.5–1.5]	X	X	X	X	X	X	Relative EV price; low = cost drop, high = stagnation
`d_tech_range` [0.5–1.5]	X	X	X	X	X	X	Relative EV range; low= overran infra, high = fast development

Legend: Policy shorthand — **I**: Infrastructure, **S**: Subsidy, **M**: Marketing, **Z**: zero-emission mandate. Policy configurations are written in {year: policy} format.

Multi-Objective Decision Analysis

To explore trade-offs and robustness, a Multi Scenario Multi-Objective Robust Decision Making (Multi Scenario MORDM) analysis is applied across all policy–uncertainty combinations. The dependent variables with outcome objectives are:

1. **Speed:** Time to reach 25% EV adoption (years, minimize),
2. **Adoption Level:** Share of EVs by 2035 (EV fleet percentage, maximize),
3. **Equity:** Inter-archetype disparity in adoption (two standard deviation across archs, minimize),
4. **Inclusion:** Zero-income household adoption rate (percentage arch 1 and 2, maximize).

Boxplots, robustness scoring, and PRIM analyses are used to evaluate how well the different per-policy combinations hold up under various futures. Thereafter, parallel coordinate plots and Pareto fronts are used to identify trade-offs and determine the most promising scenarios.

6. Results: Robust Policy Evaluation under Deep Uncertainty

This chapter presents the results of large-scale simulation experiments performed using the EMA Workbench. Building on the exploratory policy insights that informed the experimental design in Section 5.2.4, it evaluates the performance of six selected policy configurations under a wide range of future scenarios reflecting behavioural and technological uncertainties. The aim is to answer the final two sub-research questions:

- How do different future scenarios influence the speed and equity of policy interventions?
- What are the trade-offs between equity and speed objectives in relation to targeted policy interventions to promote EV adoption?

The results provide evidence on the robustness, efficiency, and inclusivity of different EV adoption strategies under deeply uncertain conditions. Based on these findings, the best overall performing policy configurations are applied to the case study city of The Hague, enabling a final analysis and answering the main research question:
What impact do targeted policy interventions have on the speed and equity of electric vehicle adoption trends in The Hague?

6.1. Scenario-Dependent Policy Performance

This section addresses SRQ2: *How do different future scenarios influence the speed and equity of policy interventions?* To evaluate this, each of the six selected policy configurations is assessed across 1,000 plausible futures generated through Latin Hypercube sampling of key behavioural and technological uncertainties.

The analysis proceeds in three steps. First, the distribution of policy outcomes across scenarios is examined to assess variability and performance stability. This helps identify which policy strategies produce consistently favourable results and which are highly context-dependent. Second, a robustness scoring method is applied, using performance thresholds across four outcome metrics to evaluate how often each policy meets critical success criteria. This enables direct comparison of the robustness of policies under uncertainty. Third, the influence of individual uncertainties on policy performance is explored to understand which factors drive variability and where interventions are most sensitive to assumptions.

Together, these analyses provide a comprehensive picture of policy reliability under deep uncertainty, setting the stage for trade-off exploration in the next section.

6.1.1. Distribution of Policy Outcomes

To assess the robustness of each policy's performance across a wide range of plausible futures, Figure 6.1 shows boxplots of the chosen four key outcome metrics across 1,000 simulated scenarios. Each boxplot visualizes the distribution of outcomes across these simulated futures per policy, which can be compared to the *No policy* control. The black horizontal line inside each box indicates the median, the value separating the better-performing half from the worse-performing half. The box spans the interquartile range (IQR), showing the middle 50% of outcomes, while the whiskers extend to 1.5 times the IQR. Dots beyond the whiskers represent outliers. This visualization captures both the typical performance and the variability or risk of each policy, highlighting not only average outcomes but also consistency and susceptibility to failure under uncertainty.

As expected, most policy combinations outperform the *No policy* case across all metrics. However, given the strength of the interventions, median differences are often modest, and wide interquartile ranges indicate that policy effectiveness varies considerably across scenarios. This is particularly evident in the top-left panel on adoption speed, where many policies still include the maximum of 14 years in their upper quartile. Only *AllPolicies* clearly reduces both the median and lower quartile, suggesting that accelerating early adoption remains challenging. Differences between policy outcomes become more pronounced in the top-right panel, which shows total EV share in 2035. Here, *Infra+Marketing* and ZEZ-related combinations, especially *Subsidy+Awareness+ZEZ* and *Infra+ZEZ*, approach the performance of *AllPolicies* and clearly outperform *No policy*. Notably, ZEZ policies raise both median and upper quartile values, whereas *Infra+Marketing* primarily improves the upper quartile. This suggests that the effectiveness of *Infra+Marketing* is more dependent on specific scenario conditions.

In the inter-archetype equity panel (bottom-left), this effect is even more pronounced. While medians are similar across policies, upper quartiles and whiskers are substantially lower under active policies, with only a few outliers exhibiting extreme inequality. The final panel on low-income inclusion reinforces this point. *EarlySubsidy+Infra* performs poorly, while *AllPolicies* and *Subsidy+Awareness+ZEZ* show higher inclusion, particularly in terms of median outcomes. Interestingly, low-income inclusion is often higher than general adoption rates, suggesting that equity-minded policies can meaningfully include lower-income groups, even when total adoption is limited.

Overall, following *AllPolicies*, the ZEZ-based combinations, particularly *Subsidy+Awareness+ZEZ*, perform best across scenarios, while *EarlySubsidy+Infra* performs consistently poorly, except on the inter-archetype equity metric. While these performance differences are meaningful, the wide interquartile ranges across all metrics raise concerns about robustness. These findings underscore the need for further robustness analysis to assess whether policies can reliably meet minimum thresholds under uncertainty. The next sections explore these aspects in more depth.

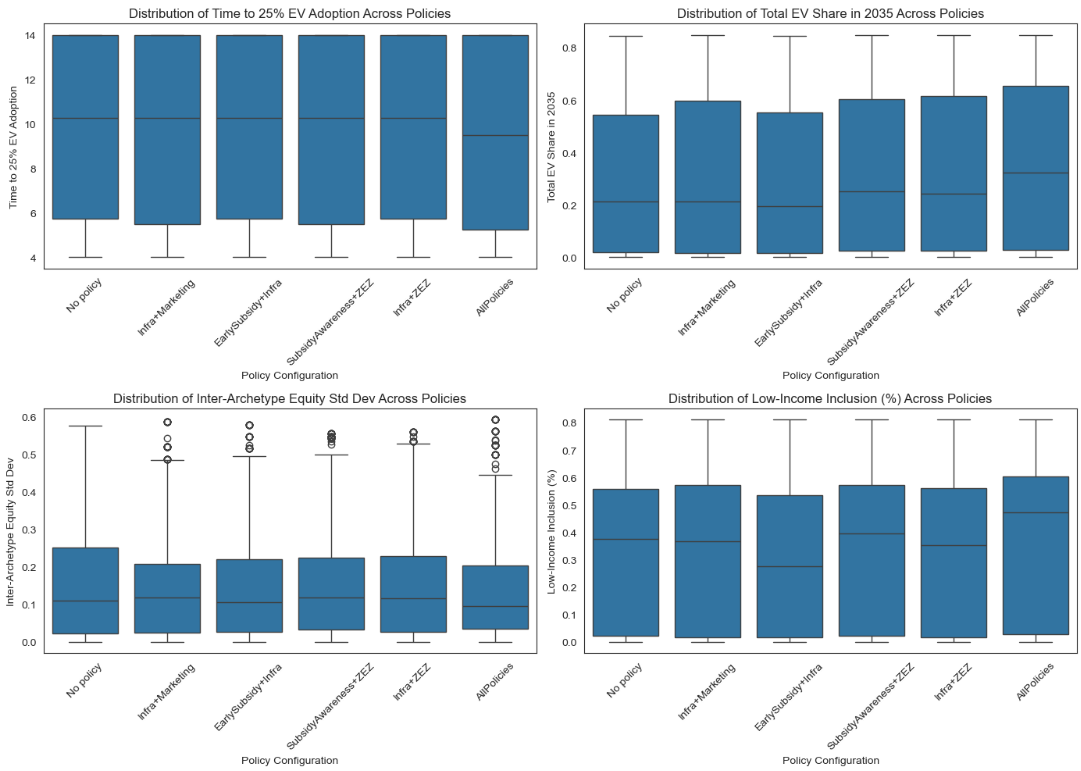

Figure 6.1.: Distribution of key outcome metrics (1,000 scenarios per policy) across six policy configurations.

6.1.2. Robustness Scoring

To evaluate how consistently each policy meets performance expectations under uncertainty, Figure 6.2 presents two complementary robustness scoring visuals. The left panel shows a stacked bar chart of the percentage of scenarios in which each policy meets 0 to 4 predefined thresholds. These thresholds correspond to four criteria: (1) adoption time $\leq$ 8 years, (2) EV share in 2035 $\geq$ 40%, (3) inter-archetype equity standard deviation $\leq$ 0.15, and (4) low-income inclusion $\geq$ 30%. The colors indicate the number of criteria met per scenario, from zero (dark blue) to all four (yellow). The right panel displays a heatmap showing the percentage of scenarios in which each individual criterion is met, disaggregated by policy. This allows us to identify which criteria are most frequently achieved and where specific policy combinations fall short.

The left panel of Figure 6.2 supports the earlier boxplot findings: it shows that all policy combinations fail to meet more than one threshold in over half of the scenarios. *EarlySubsidy+Infra* even fails to meet any of the criteria in 15% of scenarios. As for the other half, the ZEZ policies follow closely (both 22%) behind the *AllPolicies* configuration (26%) in the share of scenarios meeting all four criteria. *Infra+Marketing* also performs reasonably well, with its stack of 3 and 4 criteria combined being even slightly higher than the ZEZ options, though it scores lower on the full 4-criteria grouping. Lastly, it is notable from the left panel that scenarios where exactly two criteria are met are relatively uncommon, suggesting interdependence between certain outcomes (this aspect will be further explored in Section 6.2).

The right panel offers more detail on which exact criteria are most frequently satisfied by which combinations. Across all policies, the equity metric is most often attained, with thresholds met in over 50% of cases. In contrast, performance metrics on speed and total EV share are harder to achieve across the board, stagnating around 40%. Between policies, aside from the *AllPolicies* option, *EarlySubsidy+Infra* performs slightly better on inter-archetype equity, whereas *Subsidy+Awareness+ZEZ* performs slightly better for low-income inclusion. For the performance metrics, only *EarlySubsidy+Infra* stands out again—this time for performing relatively poorly, with scores similar to or worse than *No policy*. Interestingly, although individual thresholds are met in roughly 35–60% of cases, full satisfaction of all four is considerably rarer (as seen in the left panel), indicating misalignment in the conditions under which each criterion is fulfilled.

These results underscore the difficulty of crafting policies that are robust across a wide range of futures, and highlight the potential need for stronger interventions or strategies that reduce the plausibility of unfavourable scenario conditions. The next step, however, is to first analyse the scenarios to determine the degree of concern warranted by the observed robustness.

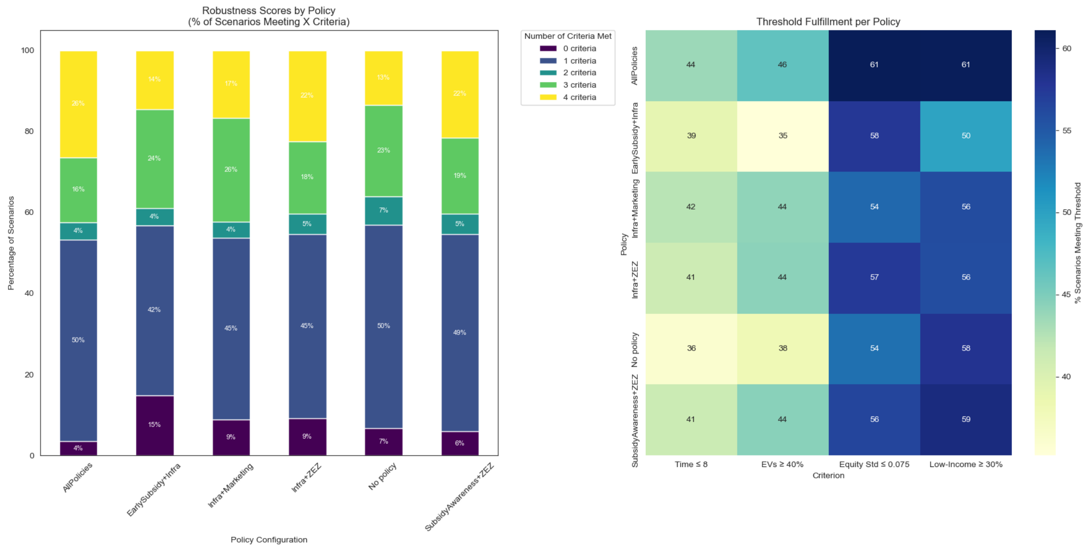

Figure 6.2.: Robustness scores. left: share of scenarios meeting 0–4 outcome thresholds. right: percentage of scenarios per policy meeting each individual threshold.

6.1.3. Scenario Discovery

Given the demonstrated influence of deep uncertainty on policy effectiveness, scenario discovery was conducted using the Patient Rule Induction Method (PRIM) to identify the conditions under which each policy configuration performs well. A scenario was considered successful using the same four criteria as in the robustness scoring of Section 6.1.2. PRIM was then applied to discover "box" regions of the uncertainty space in which the density of successful outcomes was high, while retaining sufficient coverage. Here, *density* refers to the proportion of successful scenarios within the box, while *coverage* is the proportion of all successful scenarios captured by the box.

Figure 6.3 presents the resulting trade-off plots for each policy configuration. Each point along a curve represents a PRIM box, with the y-axis showing density and the x-axis showing coverage. Color indicates the number of uncertainty dimensions restricted to define the box. These plots illustrate the inherent trade-off between generalizability and performance, where higher density often requires narrowing the box (i.e., increasing the number of restricted uncertainties). Note that this part only focuses on how well successes per policy can be predicted (i.e., the percentage in the box), and no longer on the amount of successes.

Given the six configurations tested, two main patterns arise in the PRIM trade-off plots. *EarlySubsidy+Infra*, *Infra+Marketing*, and *No policy* all show very vertical lines. They can isolate up to 80% accuracy (density), which is comparable to the other configurations and captures nearly all successful scenarios. The downside, however, is that these policy combinations require five restrictions to achieve these relatively good conditions. *Infra+Marketing* performs slightly better, reaching up to 60% accuracy with only four restrictions. Still, this again shows how not robust these configurations are, as they need substantial restriction to isolate success.

The other configurations show a more horizontal pattern after a turning point around 60%

coverage. From that point onward, they begin to lose coverage but still maintain relatively high density, with fewer restricted uncertainties. Interestingly, *Infra+ZEZ* and *Subsidy+Awareness+ZEZ* seem to offer relatively favourable trade-offs, reaching a density of 0.8 with only three restrictions.

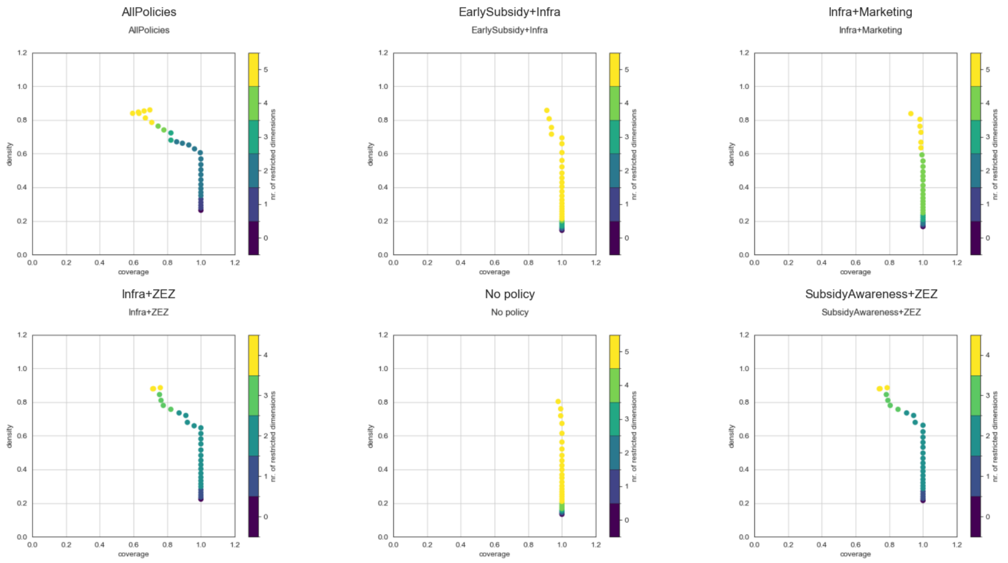

Figure 6.3.: PRIM trade-off plots per policy, showing density versus coverage when identifying high-performing subsets of the uncertainty space. Color indicates the number of restricted dimensions.

To better understand which uncertainties define the spaces of favourable outcomes, Figure 6.4 summarizes the restricted uncertainty dimensions across all PRIM boxes that achieved success. The left panel displays the median restricted ranges per uncertainty across all policy boxes that met the PRIM thresholds. In addition to the median ranges (shown as horizontal bars), minimum and maximum bounds, as well as the full tested ranges, are included to visualize the span of constrained values across policies.

The left panel shows that *d_tech_price* (EV price decline), *d_tech_range* (EV range improvement), *d_routine_buyers* (share of habitual adopters), and *d_oc_avail* (second-hand market strength) are consistently the most tightly constrained across successful boxes. Interestingly, all four variables include the current "as-is" value of 1 in their constrained ranges. This suggests that if the world unfolds according to current projections, these policies are likely to succeed. However, if conditions turn less favourable, the likelihood of success quickly drops.

d_oc_avail shows a bit more flexibility, with a median score around 0.55 on the lower side of the range. Interestingly, the whisker on the left for *d_tech_range* shows that there are outlier policy combinations that can handle less favourable conditions on that dimension. Checking the right panel, this turns out to be *AllPolicies* and also *Infra+Marketing*.

As for *d_NC_trend*, *d_routine_buyers*, and *w_pleasure*, it is clear that the successful boxes are not heavily constrained on these uncertainties—even though the literature and model highlight their importance. This could indicate that while these factors matter, they do not necessarily

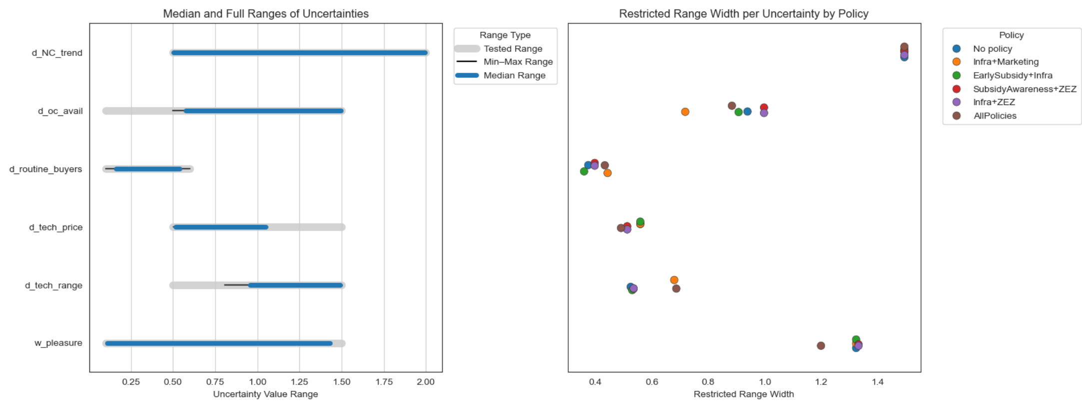

Figure 6.4.: Restricted ranges per uncertainty across all successful PRIM boxes. Median (bars) and full (black lines).

determine whether a policy is effective. Note, however, that for *d_routine_buyers*, there is still variation between policies. Specifically, *Infra+Marketing* appears somewhat sensitive to specific values, showing a smaller restricted range than the others.

Together, these figures reveal that while some policy packages are broadly robust, most rely on a favorable combination of two or three critical uncertainties. The fact that many policies share similar restricted drivers underscores the potential for targeted anticipatory action—such as promoting EV cost reductions or reducing routine buyer lock-in—to unlock success across multiple strategies. In parallel, policies that depend less on strict contextual assumptions—such as *Infra+ZEZ*—may offer more resilience in deeply uncertain environments.

Answer to RQ2

In response to SRQ2: *How do different future scenarios influence the speed and equity of policy interventions?*, the results show that different future scenarios substantially influence both the speed and equity of EV policy interventions. Some policy configurations, particularly those combining infrastructure, marketing, and ZEZs, perform well on average but rely on a favourable mix of behavioural and technological conditions to succeed. No policy is robust across all futures, though*SubsidyAwareness+ZEZ* came closest after *AllPolicies*.

Ultimately, success depends not only on policy design but also on scenario context. Factors like EV cost, driving range, and second-hand availability consistently shape outcomes. Policies that require fewer restrictive assumptions, such as *Infra+ZEZ*, may be better suited to real-world uncertainty. These findings underscore the importance of evaluating interventions not in isolation, but as part of a dynamic policy–scenario interaction space.

6.2. Trade-off Analysis Between Objectives

This section addresses SRQ3: *What are the trade-offs between equity and speed objectives in relation to targeted policy interventions to promote EV adoption?* Policymakers often face competing goals such as accelerating EV adoption while ensuring fair access across socio-economic groups. Understanding these trade-offs is critical for designing effective and inclusive transition strategies.

The analysis begins with a basic exploration of correlations to provide an initial picture of how the objectives are linked. Parallel coordinate plots are then constructed to give an overview of how different policies perform across multiple objectives and to reveal common patterns of tension or synergy between them. Pareto filtering is subsequently used to identify policies that achieve strong performance without being outperformed on any other objective, highlighting those that consistently balance competing goals. After which in the final step an significance test across all policy combinations and objectives is done to determine to which degree the found differences can be trustworthy.

Together, these methods offer a structured view of the trade space, providing insight into how trade-offs are navigated by the different policy combinations and which intervention logics yield the most balanced outcomes.

6.2.1. Correlation of Outcomes

The heatmap in Figure 6.5 visually represents the correlation matrix across policy combinations, with cell values and color intensity showing the strength and direction of relationships between the key metrics. It reveals near-perfect correlations (close to 1 or -1) between the *Time25*, *LowInc*, and *EV2035* metrics, suggesting that these parameters move together to their desired direction and are therefore not expected to necessarily trade off.

The inter-archetype equity metric based on standard deviation, however, reaches correlations of only approximately 0.3 with the performance metrics, and only 0.52 with the other equity metric (*LowInc*). This suggests two things: first, that trade-offs are expected to arise where high adoption rates mask widening gaps between neighbourhood archetypes; and second, that there is a specific difference between disparity based on archetype—which reflects education, infrastructure, and income differences—versus disparity in just the lowest-income group. In the next sections, a closer examination is provided to explore what possible trade-offs look like and how these differences manifest between the policy combinations.

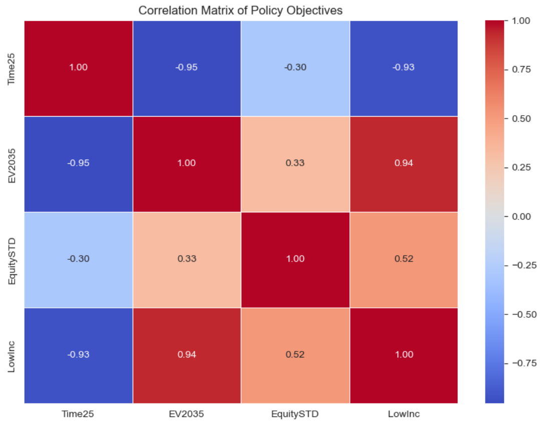

Figure 6.5.: Correlation matrix showing relationships among key outcome metrics.

6.2.2. Parallel Coordinates Plot

The parallel coordinates plot in Figure 6.6 provides a compact visualization of how each policy performs across the four normalized outcome objectives: *Speed_norm*, *EVShare_norm*, *Equity_norm*, and *LowInc_norm*. Each line represents the average normalized performance of a single policy configuration, allowing a direct comparison of trade-offs across objectives on a common scale from 0 (worst) to 1 (best). The position of each line on the axes reflects a policy's performance for that objective, with higher lines indicating better performance. The way lines intersect and space out reveals relationships between policies, where lines close together suggest similar performances and lines crossing over indicate trade-offs between objectives. The overall shape of the lines shows consistency across objectives, with straight lines indicating stable performance.

Overall, the structure of the plot reveals a consistent pattern: all policies peak sharply in the *Equity_norm* dimension while showing lower and more dispersed performance across the other three objectives. This inverted "V" shape suggests that, compared to the other metrics, *Equity_norm* is the better-performing metric, meaning it will generally perform well regardless of policy. However, the fact that the policy combinations are all so close also suggests that there is little difference in *Equity_norm* across the policies—and possibly not much difference overall (as also suggested by the boxplots in Figure 6.1). This is reinforced by the robustness analyses and could indicate that the apparent better performance in *Equity_norm*, based on normalization, may be misleading if the min and max values used are poor indicators of truly high or low performance. Therefore, trade-offs can still be interpreted, but caution is needed when interpreting the level of *Equity_norm*.

For the other metrics, differences are more apparent. *AllPolicies* again clearly stands out at the top, while *EarlySubsidy+Infra* and *No policy* perform worst. The relatively flat spread between most of the remaining lines across objectives suggests that, while no single policy is dramatically better or worse, the existing differences are potentially meaningful.

Focusing specifically on trade-offs between policies, a few clear insights emerge: *Subsidy+Awareness+ZEZ* performs better on *Speed_norm*, *EVShare_norm*, and *LowInc_norm*, but trades off on *Equity_norm* compared to the infrastructure-inclusive combinations, as indicated by its line being crossed by those in the second and third windows. Within the infrastructure policies, further variation is visible: *Infra+Marketing* performs better on *Speed_norm*, *Equity_norm*, and even slightly on *LowInc_norm*, while *Infra+ZEZ* performs somewhat better on *EVShare_norm* (as indicated by line crossings in the first and second windows).

These intersections and differences suggest that, while no single policy dominates across all dimensions, the small but consistent variations between them could still be relevant for decision-making.

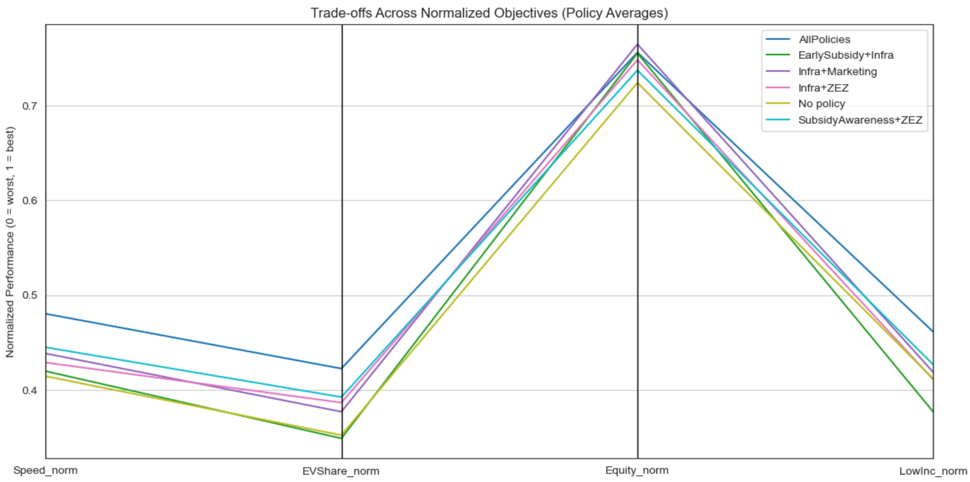

Figure 6.6.: Parallel coordinates plot showing trade-offs patterns by averages across scenarios

6.2.3. Pareto Frontier and Objective Filtering

Figure 6.7 shows a scatterplot matrix of all pairwise trade-offs between the four normalized objectives: *Speed*, *EV Share*, *Equity*, and *LowInc*. Each point represents the mean performance of a policy combination across all scenarios. Colors indicate the policy type, and shapes indicate Pareto status: triangles mark policies that are Pareto-optimal (no other policy performs at least as well on all objectives while being strictly better on one), and circles mark policies that are not Pareto-optimal because they are dominated by at least one other policy. Each subplot compares two objectives, with the upper-right corner representing high performance in both. A policy may appear competitive in one subplot but still be dominated in the full four-objective space.

Only two policies remain Pareto-optimal under this method: *AllPolicies*, which combines all measures, and *Infra+Marketing*, which targets infrastructure improvements and marketing interventions. These two policies are efficient because no other policy matches or exceeds their performance across all four objectives while improving on at least one.

The example of *SubsidyAwareness+ZEZ* (brown) illustrates why some policies are not Pareto-optimal. Although it performs reasonably well on both adoption and equity metrics, it is consistently outperformed in the full objective space. *AllPolicies* offers higher Speed and EV Share with similar or better Equity, and *Infra+Marketing* provides stronger Equity results while matching adoption performance. This means *SubsidyAwareness+ZEZ* is dominated, as there is no situation in which it would be the rational choice if all objectives are valued positively. Other non-Pareto policies fail for similar reasons, often performing well in one or two objectives but being outperformed on the remaining ones. However, *Infra+ZEZ* and *SubsidyAwareness+ZEZ* are both only dominated by *AllPolicies*. Since combining every measure is often unrealistic from a policy and resource perspective, these two can also be considered Pareto-optimal when excluding *AllPolicies* from the relevant set. Between them, *Infra+ZEZ* is unmatched in equity performance compared to *SubsidyAwareness+ZEZ*, whereas the latter is efficient due to its relatively high performance in adoption speed and EV market share. When *AllPolicies* is excluded, the rational choice set is therefore reduced to three options: *SubsidyAwareness+ZEZ*, which maximizes adoption speed and EV market share; *Infra+Marketing*, which delivers the strongest equity performance while maintaining relatively high speed; and *Infra+ZEZ*, which occupies an intermediate position, outperforming the first on equity but being outperformed by the latter on total adoption performance. Selecting any other policy would require knowingly accepting an inefficient trade-off.

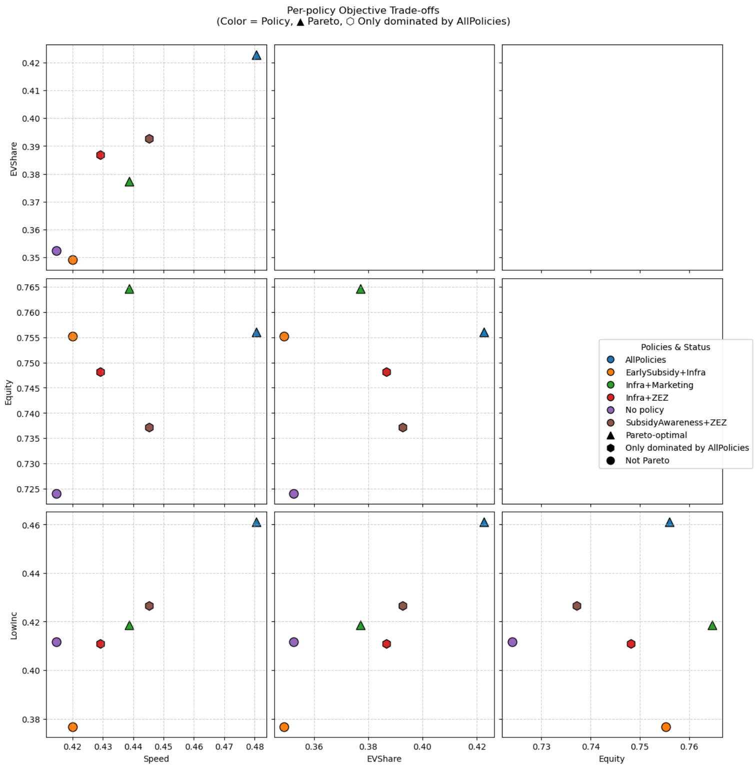

Figure 6.7.: Scatterplot matrix showing pairwise trade-offs between four normalized objectives for each policy combination, with colors indicating policy type and shapes denoting Pareto-optimality in the averaged four-objective space.

6.2.4. Statistical Significance of Policy Outcomes

To assess whether the observed performance differences between policy combinations were also relevant from the perspective of mathematical statistics, a non-parametric significance testing procedure was applied across all 1000 simulation runs per policy. As model outputs for all objectives deviated from normality (Shapiro–Wilk $p < 0.001$) and showed unequal variances (Levene's $p < 0.001$), the Kruskal–Wallis H-test was used to detect global differences, followed by Dunn's post-hoc pairwise comparisons with Bonferroni correction. This approach makes minimal distributional assumptions and accounts for multiple comparisons.

The results (full table of numbers in Appendix A) indicate that, across the four outcome metrics, the *AllPolicies* configuration consistently outperformed the baseline at conventional significance thresholds ($p < 0.05$), particularly in overall EV fleet size by 2035, time to reach 25% penetration, and low-income household adoption. Intermediate or partial combinations such as *Infra+Marketing*, *Infra+ZEZ*, and *SubsidyAwareness+ZEZ* did not produce statistically distinguishable improvements over the baseline. This suggests that, within the model structure as set in this research, only shifts both speed and equity require the harshest possible bundled intervention, while narrower packages struggle to create significant change from the baseline.

For most policies there is only limited evidence that they outperform the baseline. This reinforces the notion that there is no silver bullet. The subsequent spatial implementation analysis focuses on those policy combinations that, while not always statistically dominant, show robust and interpretable patterns in the trade-off space.

Answer to SRQ3

In response to SRQ3 *What are the trade-offs between equity and speed objectives in relation to targeted policy interventions to promote EV adoption?*, the analysis shows that while adoption speed, total uptake, and low-income inclusion are often positively correlated, these gains are not always matched by equity across neighbourhood archetypes. This indicates that high adoption can coexist with widening spatial disparities.

No single policy performs efficiently and consistently well across both scenarios and objectives. Among the better-performing ZEZ-based policies, *SubsidyAwareness+ZEZ* achieves relatively high adoption speed and market share but trades off with lower inter-archetype equity. Conversely, infrastructure-oriented policies such as *Infra+Marketing* and *Infra+ZEZ* perform more strongly on equity, while giving up some speed or total uptake. The *Early-Subsidy* option, by contrast, shows inconsistent or poor results, even underperforming the *No Policy* baseline in some trade-offs, and is therefore not recommended.

These patterns suggest that aligning equity with rapid adoption is possible, but requires integrated packages that combine infrastructure investment, targeted financial support, and measures that shift behaviour and norms. While the current tested configurations do not fully resolve the trade-off, the most balanced options emerge from mixing equity-focused measures with adoption accelerators.

6.3. Implementation in The Hague: From Strategy to Spatial Reality

This section answers the main research question of this thesis:

> *What impact do targeted policy interventions have on the speed and equity of electric vehicle adoption trends in The Hague?*

To answer this, the most promising policy configurations are applied to a spatially explicit simulation of The Hague. This final experiment bridges the gap between stylized model archetypes and a real urban setting, allowing for a realistic assessment of how policies perform when confronted with actual socio-economic diversity and spatial constraints.

Each of The Hague's 44 urban districts was simulated individually using localized socio-economic input data. This includes variation in average income, education levels, private car ownership, and housing density. The district-specific model runs use the same policy logic and behavioural rules as the main model but are initialized with empirical values, ensuring a realistic base for localized dynamics.

The selected configuration for this test, *Subsidy+Awareness+ZEZ*, combines infrastructure investment, targeted awareness campaigns, and regulatory pressure (ZEZ) in a phased manner. This policy was chosen based on its relatively robust strong performance across multiple objectives, but with a focus on speed in line with the real world needs (see Section 6.1.2 and Section 6.2). In addition, the lighter *Infra+Marketing* policy which still showed reasonably strong results and performed well on pareto optimality and interarchetype equity is therefore included as a comparative reference to assess the effectiveness of more moderate intervention strategies.

6.3.1. EV Adoption Trajectories and District-Level Outcomes

Figure 6.8 presents the simulated EV adoption trajectories for all districts from 2022 to 2038, along with a map of the final EV fleet share per district. The black dashed line in the left-hand plot shows the city-wide weighted average. Results show consistent growth across all districts, with most areas reaching between 25% and 45% EV penetration by 2035. However, substantial disparities emerge, shown by the big spread. This suggests that while the average adoption across the city is satisfactory, equity of outcomes is very uneven without spatially differentiated measures.

To better understand the variation in outcomes, the district-level results are compared to the previously defined socio-spatial archetypes(see Figure 6.9). These archetypes cluster neighbourhoods based on structural characteristics: income, education, and density. Archetype 3 (high-income, high-education, high-infrastructure), and to a lesser extent Archetype 2 (moderate income, moderate education, moderate infrastructure), generally achieve the highest EV adoption rates. In contrast, Archetypes 0 (moderate income, high education, low infrastructure) and 1 (low income, low education, low-to-moderate infrastructure)—representing groups with less favourable conditions for EV adoption according to literature —consistently underperform. Some additional variation, not fully explained by this basic archetype classification, includes the notably low adoption in the yellow area, which corresponds to the city centre of The Hague. While this may reflect a unique spatial context, the overall

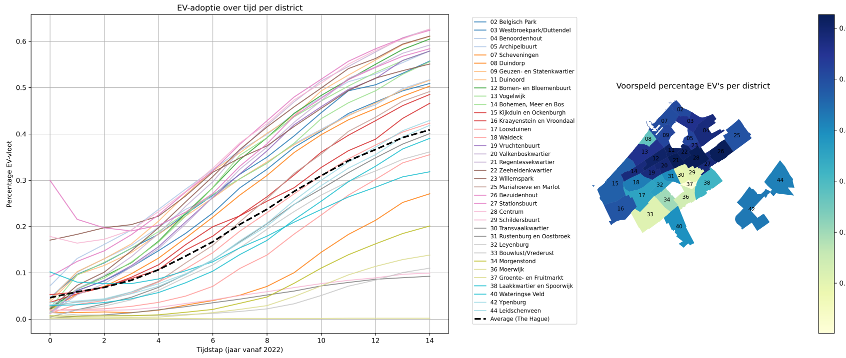

Figure 6.8.: EV adoption over time by district (left) and predicted EV fleet percentage by 2038 (right). The black dashed line represents the weighted average across The Hague.

alignment between adoption patterns and archetype categories confirms that archetypes are a useful predictor of policy responsiveness.

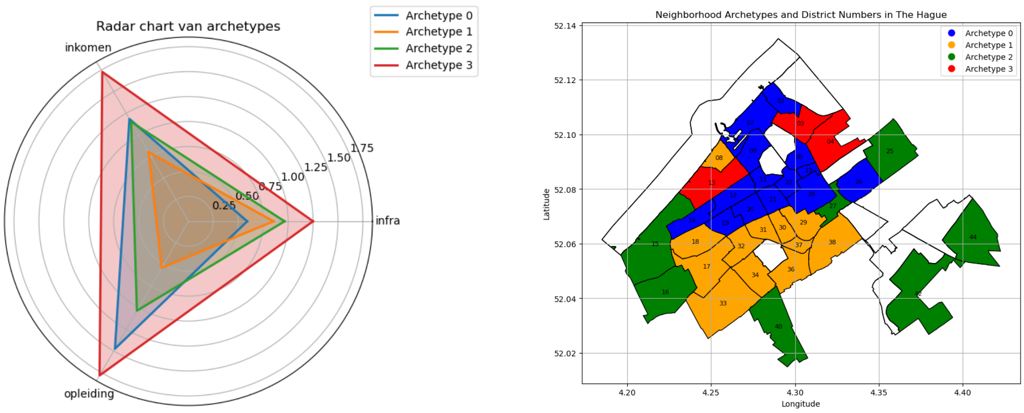

(a) (b)

Figure 6.9.: (a) Radar plot of archetype features. (b) Spatial distribution across The Hague.

6.3.2. Comparing Selected Policies in The Hague Simulation

Figure 6.10 shows the selected policies implemented in the full The Hague environment. The first point to note is that neither policy leads to a dramatic change in results. This is in line with the analyses in Section 6.1.2. Nonetheless, both policies achieve a modest improvement in mean outcomes. In particular, *Subsidy+Awareness+ZEZ* scores about five percentage points higher than the baseline on average. Notice that that differnece exceeds the number of EVs currently on the road. *Subsidy+Awareness+ZEZ* produces the highest peak adoption rates by 2038, with several districts surpassing 60% EV share and one or two outliers exceeding 70%. *Infra+Marketing* results in a smoother, more moderate acceleration, with most districts converging around the 45–55% range at the horizon. A key difference is the steeper adoption curves under *Subsidy+Awareness+ZEZ*, which is an immediate result of the zero-emission zone implementation, whereas *Infra+Marketing* delivers steadier, more gradual growth. In *Subsidy+Awareness+ZEZ*, the adoption distribution is has some districts, while lower-readiness districts lag noticeably. *Infra+Marketing* shows a slightly narrower spread, with fewer extreme leaders and laggards, suggesting it delivers more balanced growth across neighbourhood types.

Archetype 0 (blue, low infrastructure, but highly educated) responds strongly under both *Infra+Marketing* and *Subsidy+Awareness+ZEZ*, but in *Subsidy+Awareness+ZEZ* several districts in this group reach adoption levels well above the city average by mid-horizon. *Infra+Marketing* still achieves strong performance here, but differences between districts are smaller. Archetype 2 (green, almost suburban and average on all characteristics) gains more modestly in *Infra+Marketing*, with final adoption slightly above the base case; in *Subsidy+Awareness+ZEZ*, these districts improve significantly. Archetype 1 (orange, mid-level infrastructure but low education and income) climbs slowly but steadily under *Infra+Marketing*, closely following the citywide average, whereas in *Subsidy+Awareness+ZEZ* the contrast with other districts remains starker. Archetype 3 (red, high education, income and infra, often early adopters) improves over time under both policies, but *Subsidy+Awareness+ZEZ* accelerates early adoption more than *Infra+Marketing*. However, the difference compared to the base case is even smaller than for other districts, indicating that these residents might already be on track without further intervention.

Answer to the Main Research Question

What impact do targeted policy interventions have on the speed and equity of electric vehicle adoption trends in The Hague?

The spatially explicit simulation results indicate that targeted policy interventions—particularly those combining enabling measures (e.g., infrastructure investment, information campaigns) with regulatory instruments (e.g., zero-emission zones)—can slightly accelerate EV adoption city-wide. However, without additional spatial targeting, these gains remain unevenly distributed.

Comparing the two selected policies, *Subsidy+Awareness+ZEZ* delivers steeper and higher peak adoption trajectories—driven by the regulatory push of the ZEZ—but with greater disparity between leading and lagging districts. In contrast, *Infra+Marketing* produces steadier, more balanced growth, narrowing the gap between archetypes, though at the cost of slightly slower overall uptake.

These findings suggest that for maximum impact, interventions must be both robust in design and responsive to local variation in constraints and capabilities. Without spatially differentiated targeting—particularly toward low-readiness districts—policies risk reinforcing existing adoption inequalities, even when explicitly equity-minded.

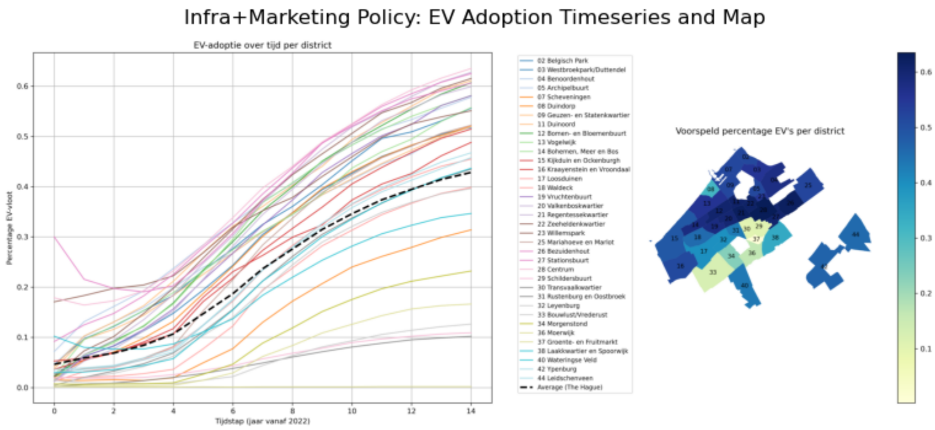

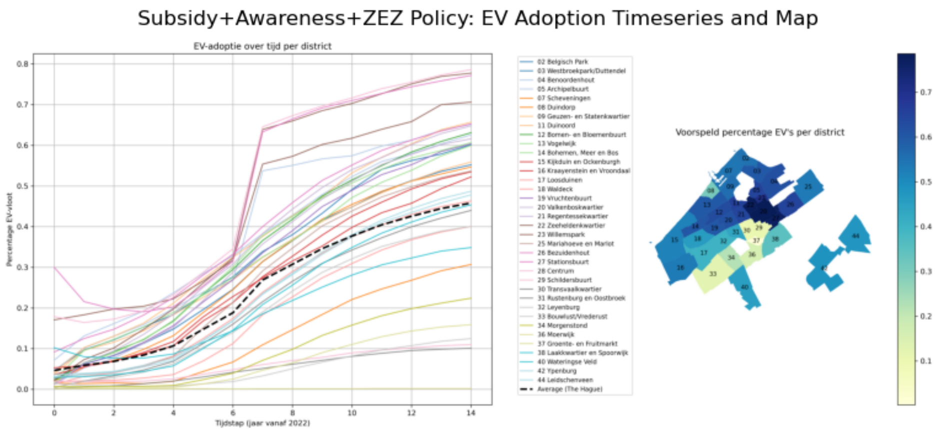

Figure 6.10.: EV adoption over time by district under *Infra+Marketing* (PC2) (upper) and under *Subsidy+Awareness+ZEZ* (PC4) (lower).

7. Discussion

7.1. The Main Findings

This study examined how different policy interventions influence the adoption of electric vehicles (EVs) in terms of both uptake speed and distributional equity.

Across 1,000 simulated futures, most targeted policy combinations outperform the *No policy* case on speed, adoption, and equity, yet the boxplots in Section 6.1.1 reveal that median gains remain modest and performance varies considerably, underscoring a strong dependence on contextual conditions. Notably, *AllPolicies* is the only configuration that consistently accelerates adoption, which highlights the difficulty of generating an immediate adoption surge. This aligns with findings by Zhang et al. [2024] and Varghese et al. [2024], which identify psychological rather than structural constraints as the primary adoption barrier, requiring time for social diffusion rather than quick resolution through infrastructure or subsidies alone.

ZEZ-based packages emerge as consistently higher-performing alternatives. Robustness scoring in Section 6.1.2 indicates that for just over 20% percent of scenarios they can achieve a 40% fleet share by 2035 and 25% by 2030, comparable to the current global leader Norway [Ritchie, 2024], while limiting inter-archetype differences to 15% and ensuring that low-income groups reach at least 30% participation in over 20% of scenarios. Although these figures may appear modest, they would represent a substantial improvement over projections by Paradies et al. [2023], which anticipate EV sales of only 40% by 2030, and over current equity gaps, with 2024 adoption differences between the lowest- and average-income groups around 30% [European Commission, 2024]. These packages combine regulatory pressure, which prompts reconsideration of vehicle choice, with complementary measures that mitigate perceived barriers of range and cost, consistent with Gao et al. [2022].

By contrast, *EarlySubsidy+Infra* performs poorly, failing to meet any threshold in 15% of scenarios and showing negligible improvement over *No policy*. This suggests that subsidies are insufficient to shift the broader population. The comparatively better performance of the lighter *Infra+Marketing* package, which approaches the results of *ZEZ*-based strategies, reinforces the view that adoption hinges more on changing perceptions than on financial incentives alone [Joshi et al., 2022]. Nevertheless, *Infra+Marketing* lacks robustness, as its effectiveness depends on favourable conditions or timely opportunities for adoption acceleration.

Scenario discovery indicates that most policies succeed under specific, relatively favourable conditions, with outcomes most strongly tied to three realizations of uncertainties: EV price decline, range improvement, and second-hand market strength. Other uncertainties, such as the *No car* trend, the share of habitual adopters, or perceived EV enjoyment, were not decisive despite suggested high importance of social identity indicators by [Li et al., 2023]. The “as-is” projections for key uncertainties often fall within favourable ranges, suggesting that, if these trajectories hold or improve, most policies—except *EarlySubsidy+Infra*—have a high likelihood of success. This suggests that policymakers can select configurations based on strategic priorities: *ZEZ*-based packages for rapid, large-scale impact, or *Infra+Marketing* for a less coercive approach, but should keep contingency measures (e.g., targeted subsidies or infrastructure expansion) prepared to counteract adverse developments in technology costs, range, or second-hand market conditions.

Trade-off analyses in Section 6.2 further refine these insights. Policies with infrastructure expansion perform slightly better on equity, measured as lower standard deviation between archetypes, whereas *Subsidy+Awareness+ZEZ* achieves stronger overall performance. Minor trade-offs also include *Infra+ZEZ* achieving marginally higher total adoption in 2035, where *Infra+Marketing* performing slightly better on equity. Pareto analysis confirms that, excluding *AllPolicies*, all three configurations are Pareto-optimal, implying that selection depends on policy priorities and resource constraints.

In the The Hague case study, neither *Subsidy+Awareness+ZEZ* nor *Infra+Marketing* radically alters the adoption landscape radically, though both deliver modest citywide gains over the baseline. *Subsidy+Awareness+ZEZ* achieves faster uptake and higher peak adoption—around five percentage points above average, with some districts surpassing 60% to 70% EV share by 2038—driven largely by the zero-emission zone. *Infra+Marketing* produces steadier, more balanced growth, with less divergence between high- and low-adoption districts. Archetype-level results show stronger responses among low-infrastructure and average-profile districts under *Subsidy+Awareness+ZEZ*, slower but steady growth in low-income and low-education areas under both policies, and limited gains for early-adopter districts. These patterns suggest that archetype-targeted deployment could enhance equity outcomes.

Overall, the findings demonstrate that effective EV adoption policy design must account for deep uncertainty and heterogeneous local conditions. Speed and equity are not inherently conflicting objectives, but achieving both requires integrated policy packages that combine financial, informational, and regulatory instruments, tailored to diverse household and neighbourhood contexts.

7.2. Discussion of Policy Implementation

While the model demonstrates the potential of various policy packages, their real-world applicability depends on institutional feasibility, social dynamics, and strategic trade-offs. This section examines three perspectives: the challenges of implementing complex interventions and the question of whether the added complexity of multifaceted packages is justified..

Translating high-performing packages into practice requires administrative capacity, political commitment, and cross-sector coordination [Howlett, 2019]. The top-performing *Subsidy+Awareness+ZEZ* strategy, which integrates financial, informational, and regulatory measures, also demands significant institutional readiness and sustained political will. Its complexity increases the risk of administrative delays, exclusion errors, and political resistance, particularly when targeted subsidies or spatially enforced ZEZs are perceived as unfair [Jordan and Matt, 2014; Minarik, 2024]. These challenges illustrate the trade-off between maximizing performance and ensuring feasibility. Simpler strategies, such as *Infra+Marketing*, may underperform somewhat on expected final adoption but are easier to implement, more adaptable, and less vulnerable to institutional bottlenecks. Adaptive policymaking frameworks stress the importance of flexible, modular measures [Haasnoot et al., 2013], suggesting that leaner options can serve as scalable entry points while more targeted elements are layered in over time. Moreover, as policy opportunities often emerge unpredictably [Hoefer, 2022; Rose et al., 2017], during a so called policy window having "good enough" options ready for rapid deployment is critical.

This could go hand in hand with the use of the suggested archetypes. The adoption patterns across archetypes closely align with the well known Rogers' diffusion of innovations theory [Rogers, 1983], with high-income, highly educated urban districts (Archetype 3) behaving as early adopters, and low-income, infrastructure-poor districts (Archetypes 0 and 1) resembling late adopters. However, adoption phases are not fixed; heterogeneity within archetypes allows for faster-than-expected uptake when local conditions—such as infrastructure or peer influence—shift favourably [Granovetter, 1978]. This study's archetype framework operationalizes Rogers' phases spatially, offering policymakers a means to target interventions more effectively by linking diffusion categories to specific neighbourhood types while accounting for intra-group variation.

This could also support cost reduction if only certain groups are eligible for the policy. It does however raise the question of which marginal gain justifies a given level of effort and budget. In particular, the modest advantage of complex packages over simpler ones prompts questions about proportionality. While *Subsidy+Awareness+ZEZ* performs best overall, its improvement over options such as *Infra+Marketing* is relatively small and, in the current model, not statistically significant in size, suggesting diminishing returns on resource investment [Howlett, 2019]. Under resource or capacity constraints, simpler packages may offer greater value, especially if well-timed and strategically targeted. Nonetheless, even seemingly modest gains—such as a the found 5-percentage-point increase in fleet share—are substantial in real terms. Given the Netherlands' current EV penetration of under 4%, such an improvement would affect hundreds of thousands of households, potentially accelerating the transition, supporting long-term climate targets, and reducing persistent inequalities. In this light, the additional effort required for more ambitious packages may still be justified despite their complexity.

7.3. Strengths and Limitations

This study presents several methodological and conceptual strengths. First, it integrates behavioural theory—specifically the CODEC framework—into a dynamic ABM, enabling a phased and realistic simulation of decision-making. Second, it incorporates socio-demographic heterogeneity in agent attributes and policy effects, embedding equity considerations from the outset rather than as a post-hoc addition. This facilitates deliberate, evidence-based policy design that proactively addresses distributional outcomes. Third, it applies robust scenario and sensitivity analysis techniques inspired by Exploratory Modelling and Analysis (EMA), providing structured insights into uncertainty and parameter relevance, and supporting informed, possibly adaptive decision-making under a range of plausible futures.

Nonetheless, several limitations must be acknowledged. First, model validation is inherently constrained due to the lack of longitudinal behavioural data on EV adoption, particularly for novel policies. Although face validation was supported by TNO mobility experts, empirical calibration remains a challenge. Second, the representation of social influence is simplified through static spatial neighbourhoods rather than dynamic or online social networks, potentially underestimating broader forms of diffusion. Third, the behavioural rules and policy implementation, especially those concerning routine buying and enablement scoring, are stylised due to limited data availability, risking oversimplification of (the elasticity of) individual preferences and context sensitivities.

Despite these constraints, the model maintains internal consistency and provides useful exploratory insight into policy dynamics under uncertainty. Future empirical data and experimental research could be used to further refine and validate key behavioural assumptions.

7.4. Future Research Directions

This research opens several avenues for further study. First, longitudinal survey or panel data could improve calibration and validation of behavioural parameters, particularly regarding the timing of decision phases and the triggers for exiting routine behaviour. Second, the model could be extended to simulate dynamic social networks, enabling more realistic representations of peer influence beyond spatial proximity.

Third, a promising direction lies in integrating EV adoption dynamics with other urban systems such as housing, energy demand, or public transport use. This would allow for multi-domain scenario analysis that captures co-benefits and trade-offs, specifically with the research now being done energy storage use potential of EVs. Additionally, exploring adaptive policy mechanisms, where governments respond dynamically to uptake trends, could help simulate more realistic governance settings, but also play into the sensitivity of impact to the deeply uncertain future.

As for the current setting, this could also be further developed. Specifically, the Equity objective, which is currently defined as the standard deviation between archetypes, could be adjusted to the max-min difference between archetypes. This would penalize inequity more strongly than the current setup. Furthermore, more extreme policy implementations and/or combinations could be explored to identify what would have been sufficient to meet the objectives. While such policies might not be very realistic, unlike the more plausible combinations used in this research, they could help provide additional perspective on the boundaries of effectiveness.

7.5. Theoretical and Practical Implications

From a theoretical standpoint, this study demonstrates that agent-based models can meaningfully incorporate structured behavioural frameworks like CODEC, bridging a persistent gap between psychological theory and formal simulation. By modelling decision phases and behavioural heterogeneity explicitly, this work contributes to the operationalisation of socio-cognitive theories within complex systems models. It also contributes to the literature on equitable sustainability transitions by demonstrating how distributional impacts can be quantified and simulated at the household level, and subsequently analysed at the neighbourhood archetype and citywide scale for policy application.

Practically, the findings provide guidance for policymakers aiming to accelerate EV adoption without reinforcing existing inequalities. The model shows that bundling targeted subsidies, spatially differentiated infrastructure investment, and information campaigns produces more inclusive outcomes than any single measure alone. Moreover, it suggests that policy effectiveness is highly sensitive to timing and behavioural context, implying once again that implementation strategies must be adaptive and evidence-informed.

The results underscore the importance of integrating equity metrics into the evaluation criteria for transport policies. By going beyond aggregate adoption rates, this study shows that transition dynamics are not uniform, and that equity considerations must be embedded in both the design and assessment phases of policy planning. Without specific attention for equity, increasing EV adoption rates will go together with increasing social inequality. This finding aligns closely with the approach to Grand Challenges—such as the energy transition—promoted by the Engineering and Policy Analysis master's program for which this thesis is written. While the answers are not conclusive, they can help decision-makers define a space of possibilities within which they can position themselves adaptively, based on the policy window, democratic needs, and supported by clear argumentation.

8. Conclusion

This thesis explored how policy interventions can influence the speed and equity of electric vehicle (EV) adoption, using an agent-based model enriched with behavioural decision-making through the CODEC framework. The findings show that while market dynamics alone may eventually lead to widespread EV uptake, they are unlikely to do so in an inclusive and timely manner. Well-timed and targeted policies—especially those combining subsidies, awareness campaigns, and regulatory tools like zero-emission zones—can accelerate adoption and reduce inequality across socio-spatial groups, but there is no silver bullet and the effects are modest.

In depth analysis demonstrates that no single policy exhibits robustness across all tested future conditions. The most influential constraining uncertainties are persistently high EV prices, insufficiently perceived driving ranges, and—less frequently identified in the literature—the underperformance of the second-hand market. Multifaceted policy packages, such as *Subsidy + Awareness + ZEZ*, achieve the highest average performance; however, their advantages over simpler strategies, including *Infra + Marketing* or *EarlySubsidy + Infra*, are moderate. A salient trade-off is observed: ZEZ-based policies outperform on conventional metrics such as adoption speed and overall uptake, whereas *Infra + Marketing* and related combinations deliver superior equity outcomes across socio-demographically diverse neighbourhoods. The projected gains from the better performing policy packages —approximately a five-percentage-point increase in EV share by 2038 relative to a no-policy baseline—appears modest but is considerable given that the current national EV share remains below four percent.

The broader discussion underscores that policy design must balance ambition with realism. Implementation dynamics, behavioural diffusion patterns, and diminishing returns all point toward the need for adaptive, phased strategies that can evolve with system readiness. Rather than treating equity and speed as trade-offs, this study shows they can be pursued jointly—but only through intentional and context-aware design.

By combining behavioural theory with spatial equity metrics, this research contributes to both the modelling literature and policy debates on sustainable mobility. It demonstrates the importance of integrating uncertainty, heterogeneity, and inclusion into transition planning, and offers a practical modelling framework for anticipating the differential impacts of EV policies across diverse urban populations.

A. Policy Significance Test

Table A.1.: Bonferroni-adjusted p-values (with significance stars) for pairwise policy comparisons across four EV adoption metrics (Dunn's test).

Row vs. Col	Pair	EV_fleet	time_25%	disparity	low_income
1	AP vs. ESI	$< 0.001^{***}$	$< 0.001^{***}$	1.000	$3.9 \cdot 10^{-9\,***}$
2	AP vs. IM	0.011^{*}	0.010^{*}	1.000	0.010^{*}
3	AP vs. IZ	0.084	0.010^{*}	1.000	0.0007^{***}
4	AP vs. NP	0.000006^{***}	0.0016^{**}	0.586	0.0009^{***}
5	AP vs. SAZ	0.403	0.284	1.000	0.078
6	ESI vs. IM	0.922	1.000	1.000	0.051
7	ESI vs. IZ	0.207	1.000	1.000	0.360
8	ESI vs. NP	1.000	1.000	0.525	0.328
9	ESI vs. SAZ	0.038^{*}	1.000	1.000	0.0064^{**}
10	IM vs. IZ	1.000	1.000	1.000	1.000
11	IM vs. NP	1.000	1.000	0.131	1.000
12	IM vs. SAZ	1.000	1.000	0.871	1.000
13	IZ vs. NP	0.317	1.000	1.000	1.000
14	IZ vs. SAZ	1.000	1.000	1.000	1.000
15	NP vs. SAZ	0.063	1.000	1.000	1.000

Note: Stars indicate Bonferroni-adjusted significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Policy abbreviations: AP = *AllPolicies*, ESI = *EarlySubsidy+Infra*, IM = *Infra+Marketing*, IZ = *Infra+ZEZ*, NP = *No policy*, SAZ = *SubsidyAwareness+ZEZ*.

B. PCA Policy Combination Clustering Results

Figure B.1 presents the clustering of top-performing policy schedules using principal component analysis (PCA) for dimensionality reduction. Each subplot corresponds to a filtered subset of scenarios: all top policies (*Top*), top policies without zero-emission zone measures (*Top (No Z)*), top policies with no more than three measures (*Top (3 Policies)*), and top policies that meet both criteria (*Top (No Z 3)*).

In each panel, points represent individual policy schedules, positioned according to their first two principal components, which capture the largest variance in the underlying performance metrics. Colours indicate cluster membership as identified by the applied clustering algorithm. The results highlight distinct groupings of policy schedules with similar performance and composition, revealing patterns in how specific measures tend to co-occur in high-performing combinations. Differences across the four panels illustrate how the removal of certain measures or limits on policy count can shift the composition and similarity structure of these top-performing policy sets

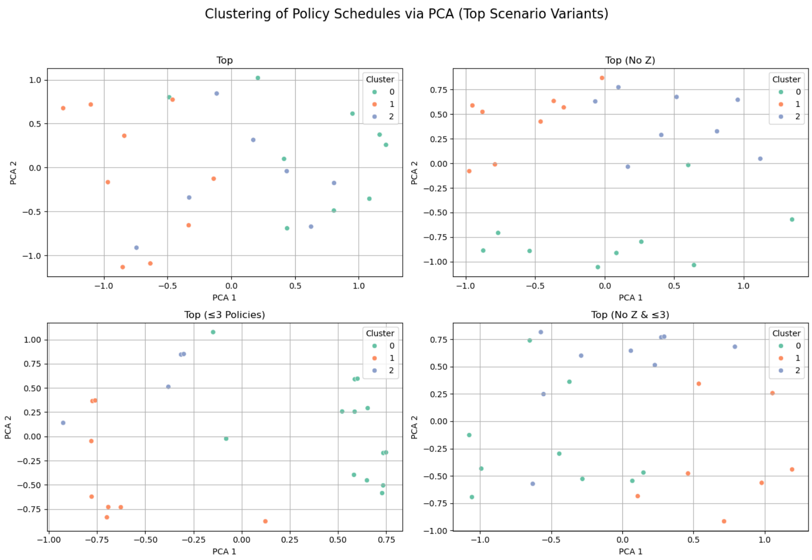

Figure B.1.: PCA clustering result

Bibliography

ANWB (2025). Alle fiscale voordelen voor elektrische auto’s op een rij — ANWB.

Axsen, J. and Kurani, K. S. (2013). Hybrid, plug-in hybrid, or electric—What do car buyers want? *Energy Policy*, 61:532–543.

Ball-Burack, A., Sun, R., Stack, S., Ou, S., Bose, R., and Yang, H.-C. (2024). Assessing the behavioral realism of energy system models in light of the consumer adoption literature. *Renewable and Sustainable Energy Reviews*, 211:115184.

Bhat, F. A. and Verma, A. (2022). A bibliometric analysis and review of adoption behaviour of electric vehicles. *Transportation in Developing Economies*, 9(1).

Buhmann, K. M., Rialp-Criado, J., and Rialp-Criado, A. (2024). Predicting Consumer Intention to Adopt Battery Electric Vehicles: Extending the Theory of Planned Behavior. *Sustainability*, 16(3):1284.

Central Bureau Statistics Netherlands CBS (2023). How much greenhouse gas does the transport sector emit?

European Commission (2024). Consumer Monitor and Survey 2023 — European Alternative Fuels Observatory.

European Commission (2025). Consequences of climate change.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7(2):117–140.

Franke, T., Neumann, I., Bühler, F., Cocron, P., and Krems, J. F. (2011). Experiencing Range in an Electric Vehicle: Understanding Psychological Barriers. *Applied Psychology*, 61(3):368–391.

Gao, Y., Leng, M., Zhang, Y., and Liang, L. (2022). Incentivizing the adoption of electric vehicles in city logistics: Pricing, driving range, and usage decisions under time window policies. *International Journal of Production Economics*, 245:108406.

Gemeente Den Haag (2025). Milieuzones Den Haag - Den Haag.

Gnann, W., Ziegler, A. F., and Tietge, U. (2023). Inequitable access to ev charging infrastructure. *Environmental Science Technology*, 57(2):1120–1128.

Granovetter, M. (1978). Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., and Ter Maat, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change*, 23(2):485–498.

Hertzke, P., Schaufuss, P., Kampshoff, P., and Möller, T. (2025). New twists in the electric-vehicle transition: A consumer perspective.

Hoefer, R. (2022). The Multiple Streams Framework: Understanding and Applying the Problems, Policies, and Politics Approach. *Journal of Policy Practice and Research*, 3(1):1–5.

Hopkins, E., Potoglou, D., Orford, S., and Cipcigan, L. (2023). Can the equitable roll out of electric vehicle charging infrastructure be achieved? *Renewable and Sustainable Energy Reviews*, 182:113398.

Howlett, M. (2019). *The policy design primer*.

Huang, X., Lin, Y., Zhou, F., Lim, M. K., and Chen, S. (2021). Agent-based modelling for market acceptance of electric vehicles: Evidence from China. *Sustainable Production and Consumption*, 28:206–217.

IEA (2024). Global EV Outlook 2024 – Analysis - IEA.

Jain, M., Talwar, S., Rastogi, R., Kaur, P., and Dhir, A. (2024). Policy stimulation for the electric vehicle industry: An analysis of mainstream media discourse. *Business Strategy and the Environment*, 33(6):5303–5324.

Jordan, A. and Matt, E. (2014). Designing policies that intentionally stick: policy feedback in a changing climate. *Policy Sciences*, 47(3):227–247.

Joshi, N., Malhotra, M., and Singh, J. (2022). Assessing adoption intention of electric vehicles in India: The mediating role of government policies. *Deleted Journal*, 22(1).

Kaas, Bart, B., Causevic, S., and Public, T. (2024). Definition of archetypical neighborhoods for residential flexibility analyses. Technical Report TNO 2024 R11524.

Lee, H. J. (2020). A Study of Consumer Repurchase Behaviors of Smartphones Using Artificial Neural Network. *Information*, 11(9):400.

Lee, Y.-I., Vu, A., and Trim, P. (2021). Millennials and repurchasing behaviour: a collectivist emerging market. *International Journal of Retail Distribution Management*, 50(5):561–580.

Li, W., Wang, M., Cheng, X., and Long, R. (2023). The impact of interaction on the adoption of electric vehicles: Mediating role of experience value. *Frontiers in Psychology*, 14:1129752.

Liao, F., Molin, E., Timmermans, H., and Van Wee, B. (2018). Carsharing: the impact of system characteristics on its potential to replace private car trips and reduce car ownership. *Transportation*, 47(2):935–970.

Mandolakani, F. S. and Singleton, P. A. (2024). Electric vehicle charging infrastructure deployment: A discussion of equity and justice theories and accessibility measurement. *Transportation Research Interdisciplinary Perspectives*, 24:101072.

Martens, K. and Golub, A. (2018). A Fair Distribution of Accessibility: Interpreting Civil Rights Regulations for Regional Transportation Plans. *Journal of Planning Education and Research*, 41(4):425–444.

Mashhoodi, B. and Van Der Blij, N. (2020). Drivers' range anxiety and cost of new EV chargers in Amsterdam: a scenario-based optimization approach. *Annals of GIS*, 27(1):87–98.

McDonald, J., McDonald, S., and McDonald, D. (2022). Financing electric vehicle adoption: Barriers and opportunities. *Energy Policy*, 159:112–123.

Minarik, J. J. (2024). Targeting public spending: Means-testing and user charging. *OECD Journal on Budgeting*, 23(3).

Ministerie van Algemene Zaken (2019). Klimaatakkoord.

More, Scott Hardman3nbsp;Kelly L. Flemingnbsp; andnbsp;Eesha Kharenbsp;, A. (2025). A perspective on equity in the transition to electric vehicles.

Mulley, D., Nelson, J., and Hensher, D. (2020). The role of urban form in influencing travel behavior: A review of the literature. *Transport Reviews*, 40(3):335–356.

Municipality The Hague (2024). Den Haag in cijfers.

Nikolic, I. and Ghorbani, A. (2011). A method for developing agent-based models of socio-technical systems. *International Conference on Networking, Sensing and Control*, pages 44–49.

Novizayanti, D., Prasetio, E. A., Siallagan, M., and Santosa, S. P. (2021). Agent-Based Modeling Framework for Electric Vehicle Adoption Transition in Indonesia. *World Electric Vehicle Journal*, 12(2):73.

Paradies, G. L., Usmani, O. A., Lamboo, S., and Van Den Brink, R. W. (2023). Falling short in 2030: Simulating battery-electric vehicle adoption behaviour in the Netherlands. *Energy Research Social Science*, 97:102968.

Pasaoglu, G., Harrison, G., Jones, L., Hill, A., Beaudet, A., and Thiel, C. (2015). A system dynamics based market agent model simulating future powertrain technology transition: Scenarios in the EU light duty vehicle road transport sector. *Technological Forecasting and Social Change*, 104:133–146.

Pei, M., Huang, Z., Zhang, Z., Wang, K., and Ye, X. (2025). Range anxiety and willingness to pay: Psychological insights for electric vehicle. *Journal of Renewable and Sustainable Energy*, 17(1).

Quaglieri, L., Mercuri, F., and Fraccascia, L. (2024). Investigating Consumer Behaviour towards Electric Vehicles: A Systematic Literature review. *Circular Economy and Sustainability*.

Querini, F. and Benetto, E. (2014). Agent-based modelling for assessing hybrid and electric cars deployment policies in Luxembourg and Lorraine. *Transportation Research Part A Policy and Practice*, 70:149–161.

Ritchie, H. (2024). Tracking global data on electric vehicles.

Rogers, E. M. (1983). Diffusion of Innovations. *SSRN Electronic Journal*.

Rose, D. C., Mukherjee, N., Simmons, B. I., Tew, E. R., Robertson, R. J., Vadrot, A. B., Doubleday, R., and Sutherland, W. J. (2017). Policy windows for the environment: Tips for improving the uptake of scientific knowledge. *Environmental Science Policy*, 113:47–54.

Rudolph, C. (2016). How may incentives for electric cars affect purchase decisions? *Transport Policy*, 52:113–120.

Siebenhofer, M., Ajanovic, A., and Haas, R. (2021). How Policies Affect the Dissemination of Electric Passenger Cars Worldwide. *Energies*, 14(8):2093.

Sierzchula, W., Bakker, S., Maat, K., and Van Wee, B. (2014). The influence of financial incentives and other socio-economic factors on electric vehicle adoption. *Energy Policy*, 68:183–194.

Sikder, S. K., Nagarajan, M., and Mustafee, N. (2023). Augmenting ev charging infrastructure towards transformative sustainable cities: An equity-based approach. *Technological Forecasting and Social Change*, 196:122829.

Soltani Mandolakani, F. and Singleton, P. A. (2024). Electric vehicle charging infrastructure deployment: A discussion of equity and justice theories and accessibility measurement. *Transportation Research Interdisciplinary Perspectives*, 24:101072.

Song, R. and Potoglou, D. (2020). Are Existing Battery Electric Vehicles Adoption Studies Able to Inform Policy? A Review for Policymakers. *Sustainability*, 12(16):6494.

TNO, Sustainable Urban Mobility Safety (2023). EV charging locations in NL.

United Nations (2025). Causes and effects of climate change — United Nations.

Varghese, A. M., Menon, N., and Ermagun, A. (2024). Equitable distribution of electric vehicle charging infrastructure: A systematic review. *Renewable and Sustainable Energy Reviews*, 206:114825.

Wang, J., Huang, C., He, D., and Tu, R. (2023). Range Anxiety among Battery Electric Vehicle Users: Both Distance and Waiting Time Matter. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 67(1):1309–1315.

Zhang, L., Van Lierop, D., and Ettema, D. (2024). Electrifying: What Factors Drive the Transition Toward Electric Vehicle Adoption in the Netherlands? *Transport Policy*.

Colophon

This document was typeset using LaTeX, employing the KOMA-Script class `scrbook`, with Palatino as the primary typeface. For select non-privacy-sensitive components of the writing, structuring, and proofreading process, support was provided by OpenAI's ChatGPT, which was utilized as a writing and coding assistant after original text, ideas and sources were given in the prompt.