

ADVERTIMENT. L'accés als continguts d'aquesta tesi queda condicionat a l'acceptació de les condicions d'ús estableties per la següent llicència Creative Commons:  <https://creativecommons.org/licenses/?lang=ca>

ADVERTENCIA. El acceso a los contenidos de esta tesis queda condicionado a la aceptación de las condiciones de uso establecidas por la siguiente licencia Creative Commons:  <https://creativecommons.org/licenses/?lang=es>

WARNING. The access to the contents of this doctoral thesis is limited to the acceptance of the use conditions set by the following Creative Commons license:  <https://creativecommons.org/licenses/?lang=en>



Università
Ca' Foscari
Venezia

Modelling the social, cultural and technological evolution of low-carbon lifestyles

A dissertation submitted for the joint degree of
Doctor of Philosophy in Environmental Science and Technology
and
Doctor of Philosophy in Economics

by
Daniel Torren Peraire

Supervisors:
Prof. Jeroen van den Bergh
Prof. Ivan Savin
Prof. Enrica De Cian

Institut de Ciència i Tecnologia Ambientals,
Universitat Autònoma de Barcelona
Department of Economics,
Università Ca' Foscari Venezia

May 2025

Acknowledgements

I would like to thank Prof. Jeroen van den Bergh for his guidance, support and supervision throughout the PhD trajectory. Especially for the multitude of exchanges of ideas, often at opposite ends of the argument, we had over the last four years. I would like to extend identical thanks to Prof. Ivan Savin for his co-supervision. His precise feedback and insights on my work made the PhD journey significantly smoother. Additionally, I would like to thank Prof. Enrica de Cian for providing me with a welcoming environment during my time at Ca'Foscari.

To my two co-authors, Miquel and Tomasso. Miquel, clearly the 16th version of the model is the charm, I never want to have another google meet again, thank you for your excellent ideas and putting up with my inability to know how long things take to do. Tommaso, you're an excellent researcher, discussant, flat mate, gym buddy and as of our paper ABMer too, looking forward to more cycling trips in the future! Thank you also to Diana, Ulrike and Herbert for all their hard work in organising and running the EPOC consortium.

To all those research group members and office mates I have had the pleasure of sharing my time with Jozef, Oskar, Pablo, James, Elisa, Hugo, Austin, Enrico, Ioanna, Sofia, Valentina and the rest of the amazing post-docs thanks for helping me out when I was abysmally confused or had a crisis of research faith. Als icters i boulderers, amb menció especial a la Berta, el Mateo, el Miquel i la Maria, heu sigut la raó pels meus moments millors i sigut una ajuda constant durant els pitjors, gràcies a totes i tots, ens veiem a cycling o al rocòdrom. On a wider note, the support and joy provided by the Icta community is a large part as to why I have been able to finish. To Mark, Rob and Ben thanks for being for a source of new ideas and hilarity even from half a world away. To Alfie, Tom, Jack (the shed will forever be my home), Hunain, Alex, Dhruv, Francis, Johnny, Amy, Ollie, Imy, Freya and Ophie thanks for hosting and welcoming me back in London, love you all.

Per la part més important, gràcies a tota la família, no sabeu el suport emocional que m'heu donat durant els últims anys i continueu donar-me. Per la iaia, l'avi, la tieta, el tiet la Júlia i el Miquel, una de les raons que vaig escollir fer aquest doctorat en particular va ser per connectar millor amb vosaltres. M'heu donat consells, menjar i suport, us deu un deute enorme per això. A la baba i l'avi Daniel, gràcies per deixar-me gaudir de la vostra companyia. A la Maria i (Dr?) Paula m'heu ensenyat mil coses, i m'heu donat mil il·lusions, us estimo profundament i no sabeu l'alegria que em fa créixer amb vostres. A la mama i el papa, gràcies per sempre està amb mi, sempre confiar en mi, igualment de què lluny que he estat sempre sabia que podia comptar amb vosaltres.

This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 956107, "Economic Policy in Complex Environments (EPOC)".

Contents

1	Introduction	4
1.1	Background	4
1.2	Research methods and objectives	6
1.3	Thesis Outline	7
2	An agent-based model of cultural change for a low-carbon transition	12
2.1	Introduction	12
2.2	Cultural change and identity	13
2.3	Model components	15
2.3.1	General structure of the model	15
2.3.2	A dynamic model of culture using environmental identity	17
2.3.3	Information diffusion through imperfect imitation	18
2.3.4	Homophily and asymmetric weightings in a social network	19
2.3.5	Baseline experimental set-up	21
2.4	Results	23
2.4.1	Environmental identity dynamics	23
2.4.2	Impact of model components	25
2.4.3	Attitude change through green influencers	27
2.4.4	Sensitivity analysis	30
2.5	Limitations and outlook	31
2.6	Conclusions	33
Appendices		40
2.A	Additional figures	40
3	The cultural multiplier of climate policy	43
3.1	Introduction	43
3.2	The model	44
3.2.1	Conceptual approach	44
3.2.2	Market module	45
3.2.3	Social imitation module	46
3.2.4	Climate policy module	49
3.3	Results	49
3.3.1	Overview of experimental runs	49
3.3.2	Evaluating the cultural multiplier	52
3.3.3	Impact of key parameters on the cultural multiplier	54
3.4	Conclusions	58

Appendices	62
3.A Analytical results for the NCES utility function	62
3.B Imitation of consumption	65
3.C Additional simulation results	67
3.C.1 Network structure and homophily in initial environmental identity	70
4 Driving in the wrong direction? A co-evolutionary model of EV adoption and innovation*	73
4.1 Introduction	73
4.2 The model	74
4.2.1 Overview	74
4.2.2 Car users	75
4.2.3 Car supply	77
4.3 Calibration and validation	79
4.4 Policy experiments	80
4.4.1 Experimental setup	80
4.4.2 Policy results	83
4.5 Conclusions	90
Appendices	98
4.A Model equations	98
4.A.1 Car users	98
4.A.2 Car supply	100
4.B Model parameters and variables	103
4.C Parameter calibration	109
4.C.1 Vehicle choice parameters	109
4.C.2 Social network parameters	111
4.C.3 Exogenous car parameters	112
4.C.4 NK landscape parameters	113
4.C.5 Car manufacturer parameters	114
4.C.6 Event-based probabilities	114
4.D Sensitivity analysis	114
4.E The technology landscape	116
4.E.1 How NK landscapes work	116
4.E.2 NK landscape properties	117
4.F Lifecycle evaluation derivation	118
4.G Optimal price derivation	118
4.H Additional figures	120
5 Conclusions	125

Chapter 1

Introduction

1.1 Background

While the remaining planetary carbon budget for 1.5°C or 2°C is rapidly running out, implemented climate policy solutions are still too incomplete or weak to deliver a transition to a low-carbon economy. A challenge for policy aimed at addressing climate change is the complex adaptive nature of the economy (Arthur, 1994; Balint et al., 2017; Arthur, 2021). The complexity emerges from the interaction of many actors with diverse priorities and characteristics. Moreover, agents' abilities to adopt low-carbon alternatives are heterogeneous, due to differing needs, circumstances and economic inequality. Therefore, tailored incentives may be required to avoid current lock-in of high-carbon consumption habits and technology. On the other hand, the economy's adaptive nature arises from the fact that interactions among actors are not externally imposed. Instead, these exchanges, such as social imitation among consumers or firms responding to changing preferences, evolve endogenously within the system. These two characteristics combined result in the dynamic behaviour of agents that are highly non-linear and out of equilibrium. Policy-makers attempting to guide such an economy to achieve climate goals must be aware of these properties when considering policy tools.

Meeting climate goals requires radical changes in the consumption behaviour of individuals (Wynes and Nicholas, 2017). This necessitates an understanding of how the diffusion of low-carbon behaviour will occur. This thesis aims to study climate policies using models that capture key elements of the complex adaptive economy, such as the role of social interactions in consumption choices. Low-carbon behaviours do not change in isolation. Instead, they are subject to social influence through imitation or adherence to social norms. When simulating the role of social influence, opinion dynamics models offer valuable insights into how populations may converge toward consensus or fragment into polarized groups (Deffuant et al., 2002), as well as how transitions between these states unfold (Helfmann et al., 2021; Castellano et al., 2000).

Models of pro-environmental behavioural diffusion, such as energy conservation or recycling, offer insights into the key drivers and obstacles influencing the adoption of low-carbon lifestyles (Ajzen, 1991; Niamir, Kiesewetter, et al., 2020). Shifts often involve interconnected choices across different behaviours, such as following a vegetarian diet, purchasing an electric vehicle, or avoiding air travel. Therefore, a multi-dimensional perspective is essential due to behavioural spillovers, where adopting one low-carbon action can influence the likelihood of adopting others such as not owning a car and following a vegan diet (Lacasse, 2016; Andersson and Nässén, 2022).

Definitions of culture are almost as diverse as the cultural variation they attempt to describe. For instance, Durham (1991) defines culture as a “system of symbolically encoded conceptual phenomena that are socially and historically transmitted within and between populations”. A common thread in many definitions is the emphasis on social learning, defined as knowledge acquisition through the involvement of others, such as imitation (Mesoudi, 2016). Cultural evolution theory describes cultural change as a Darwinian evolutionary process, based on the variation, fitness and inheritance of diverse cultural traits (Boyd and Richerson, 1988). Approaches grounded in this theory can provide an understanding of how a cultural trait spreads in a population, leading to behavioural change (Buenstorf and Cordes, 2008; Davis et al., 2018; Kaaronen and Strelkovskii, 2020). Alternatively, non-evolutionary models like the one by Epstein and Axtell (1996) represent culture as a fixed-length binary ‘cultural tag’ that encodes multiple traits. This approach captures a multi-dimensional view of culture, where agents carry several attributes shaped by social interaction. Such modelling frameworks of cultural change can identify pathways to avoid current lock-in towards high-carbon norms and shallow behavioural shifts that would limit decarbonisation (Davis et al., 2018).

Given the need for behavioural change at a systemic level, policy instruments that align the incentives of both individuals and firms with climate goals play a critical role in enabling a low-carbon transition. One of the most prominent examples is carbon pricing, either through a carbon tax (as in Sweden and the UK) or a cap-and-trade system (as the EU-ETS). Such a policy aims to internalise part of the external costs of carbon emissions by pricing them. Polluters then face a cost trade-off: pay the carbon price or invest in decarbonising their production.

For deep decarbonisation, the design of carbon pricing needs to account for consumption choices being affected not only by pricing but also by social influence (Wang et al., 2021), as the policy’s goal is to reduce emissions through behavioural change. This influence can be modelled via endogenous preferences. One way to frame this effect is through the concept of a “social multiplier” (Glaeser et al., 2003; Konc et al., 2021), in which individuals misinterpret carbon price-induced changes in others’ consumption as shifts in underlying preferences. The effect results in greater decarbonisation for the same carbon price when social interactions are considered. Current studies modelling endogenous preferences in the context of climate policy have focused on single consumption decisions (Konc et al., 2021; Mattauch et al., 2022). However, stringent carbon pricing may trigger changes across a larger number of consumption choices as households adjust from high to low-carbon alternatives across areas such as transport or food. Understanding these systemic shifts requires a broader cultural perspective (Schumacher, 2015; Bezin, 2019) on how social influence operates simultaneously across interconnected consumption behaviours. This lens can inspire a parallel concept to the “social multiplier” framed as a “cultural multiplier”, which captures how social learning across multiple consumption categories may amplify lifestyle changes via longer-term shifts in norms and environmental identity.

The transport sector is a clear example of how individual behaviour, technological lock-in, and policy interact within a complex adaptive system. Shaped by historical mobility patterns and car-centric infrastructure development, the sector demonstrates how difficult it is to shift entrenched habits without coordinated interventions. As a result, decarbonisation remains limited, with road transport alone accounted for 70% of global transport emissions in 2019 (Jaramillo et al., 2022, p 1052). EVs play a critical role in addressing this challenge, yet their diffusion depends on the co-evolution of consumer

behaviour and firm innovation. For example, “first-movers” are eager to adopt this new alternative to the internal combustion engine while more conservative “laggards” stifle the social contagion (Rogers, 2003). This may hamper investment by firms (Windrum et al., 2009) required for a complete decarbonisation of light transport. Another reason for a slow transition is that policymakers must balance the need to decarbonise with multiple other social or political objectives, such as reducing inequity or achieving a balanced budget. For example, adoption subsidies tend to benefit well-off households while negatively affecting the public budget (Herrmann and Savin, 2017; Nuñez-Jimenez et al., 2022). Therefore, policy mixes that leverage social dynamics while balancing equity, cost-effectiveness, and environmental outcomes are crucial in catalysing a transition to low-carbon mobility.

1.2 Research methods and objectives

This thesis uses agent-based modelling (ABM) as its primary method. ABMs are particularly well suited to capture the adaptive complexity of the economy, given that it allows for heterogeneity and bounded rationality of agents, social influence over networks, and dynamics of behaviour, markets and technologies (Savin et al., 2023). For example it has been used to address a wide range of topics, such as misinformation spread (Di Francesco and Torren-Peraire, 2024), pro-environmental behavioural change (Niamir, Filatova, et al., 2018; Kaaronen and Strelkovskii, 2020), energy policy (Rai and Robinson, 2015; Castro et al., 2020), electric vehicle (EV) adoption (Eppstein et al., 2011; Silvia and Krause, 2016) and macro-ABMs that aim to give a detailed representation of the economy as a whole (Lamperti et al., 2018; Safarzyńska and Bergh, 2022). ABMs follow a bottom-up approach where the behaviour of each agent (individuals, firms, institutions etc) are simulated separately. Agents are boundedly rational, may have heterogeneous characteristics and follow behavioural rules that determine how they relate to with one another and respond to their environment. Through these exchanges, ABMs reproduce complex behaviour from local microeconomic interactions.

The philosophy behind the ABM approach may be summarised in two key quotes. The first from Read, 1914 that ”it is better to be vaguely right than exactly wrong” emphasises that while non-ABM closed form models may be more elegant to work with, this often comes at the cost of making simplifications in the name of mathematical tractability. The second ”If you didn’t grow it, you didn’t explain it” (Epstein, 1998) centres the role of generating complexity behaviour endogenously to make sure that one understands the precise causal implications of interactions among many agents. This is especially pertinent when considering the complexity of social and technological processes that part of a transition to a low-carbon economy.

ABMs use different network structures to emphasise particular social interactions. For example, the Watts-Strogatz graph has a small-world property combining high clustering with short mean path length, which can describe real-world social exchanges between individuals who are each other’s physical neighbours. Alternatively, graphs with high heterogeneity in the number of connections between individuals, such as the Barabasi-Albert scale free graph, can well represent communication on social media, where a small clique has much larger influence.

This thesis uses ABMs to address the following research questions:

- How does pro-environmental diffusion of behaviour interact with cultural evolution

on a path to a low-carbon economy? To answer this, I explore the longer-term impact of cultural change and the mechanisms behind behavioural decarbonisation.

- How does cultural change moderate the effectiveness of carbon taxes. In pursuit of this goal, I analyse the magnitude of the cultural multiplier compared to the social multiplier and what socioeconomic characteristics affect this magnitude, such as network structure and substitutability of goods.
- What policy combinations that meet EV adoption targets can best balance the trade-offs between costs, emissions, and consumer utility? To address this, I first examine the policy stringency required for individual instruments to achieve a 95% electric vehicle fleet in California by 2035 and then expand this analysis to policy mixes.

1.3 Thesis Outline

The remainder of the thesis contains three studies to address the research questions. Taken together, these chapters explore climate policy through the lens of the economy as a complex adaptive system, emphasizing how individual behaviours, cultural dynamics, and technological evolution influence one another in shaping decarbonisation outcomes.

Chapter 2 develops an ABM to study how behavioural decarbonisation interacts with longer-term cultural change. Culture is defined here as socially transmitted information that shapes individual preferences and behaviours over time. Within the model, individuals make multiple lifestyle choices, such as not flying, avoiding car ownership, or adopting a vegan diet, which evolve through imperfect social learning in a network. To capture cultural change, the model focuses on environmental identity. This identity is represented as an average of attitudes toward relevant lifestyle choices and is treated as independent of behavioural outcomes. This allows for contradictions between what individuals believe and how they act. Furthermore, the strength of social influence between individuals is determined by the similarity in their environmental identity, leading to dependencies between behaviours and spillovers in pro-environmental attitudes. Lastly, green influencers are introduced as a minority of individuals who broadcast a strong pro-environmental attitude. This scenario explores how cultural dynamics can support a green transition beyond what would emerge through social diffusion alone.

To achieve emission reduction goals, the design of climate policy must account for consumption choices being influenced not only by pricing but also by social interaction. This involves extending the notion of “social multiplier” of climate policy to the context of multiple consumer needs while allowing for behavioural spillovers between these, giving rise to a “cultural multiplier”. Chapter 3 builds on the ABM introduced in Chapter 2, to assess how carbon pricing affects cultural change and, consequently, decarbonisation. In the model individuals make consumption choices between low- and high-carbon goods across multiple categories. These individuals have heterogeneous and dynamic preferences for low-carbon goods, which evolve through repeated and weighted social exchanges. To test different mediums of social influence three graphs are used in the model. Overall, the model assesses how this cultural multiplier contributes to the effectiveness of carbon taxation.

In Chapter 4, a third study extends the adaptive complexity perspective to technological innovation, maintaining a focus on the role of social imitation in networks as in

Chapters 2 and 3. To this end, the chapter develops a novel ABM of adoption of electric vehicles to examine the interaction of consumer behaviour, firm innovation, and policy incentives. This is applied to the case of California for the period 2001 to 2023. This chapter builds on ABMs of EV diffusion by bringing together literature strands that focus on the role of opinion dynamics (McCoy and Lyons, 2014; Silvia and Krause, 2016) and endogenous innovation(Windrum et al., 2009; Greene et al., 2014; Fan and Li, 2025). In the model, heterogeneous vehicle users, influenced by the purchasing behaviour of their peers, choose between new and used cars based on discrete choice theory. At the same time, manufacturers develop new vintages of EVs and conventional cars and adjust prices accordingly. Vehicle innovation occurs by exploring a multi-dimensional NK-landscape (Kauffman, 1993), thus accounting for the path-dependent nature of innovation. Specifically, firms improve their cars over four characteristics: production cost, fuel efficiency, battery or fuel tank size, and overall quality measure. Different policy combinations are examined to achieve deployment targets while balancing economic costs, emissions and consumer utility. These include carbon pricing, new and used car purchase rebates, production subsidies and electricity price subsidies. The analysis compares individual and combined policy instruments to see in what way the best results can be achieved.

A final chapter summarises the results, draws general conclusions, and offers suggestions for further research.

References

Ajzen, Icek (1991). “The theory of planned behavior”. In: *Organizational behavior and human decision processes* 50.2, pp. 179–211.

Andersson, David and Jonas Nässén (2022). “Measuring the direct and indirect effects of low-carbon lifestyles using consumption data”. In: *Journal of Cleaner Production*, p. 135739.

Arthur, W Brian (1994). “Inductive reasoning and bounded rationality”. In: *The American economic review* 84.2, pp. 406–411.

Arthur, W Brian (2021). “Foundations of complexity economics”. In: *Nature Reviews Physics* 3.2, pp. 136–145.

Balint, Tomas, Francesco Lamperti, Antoine Mandel, Mauro Napoletano, Andrea Rovennini, and Alessandro Sapio (2017). “Complexity and the economics of climate change: a survey and a look forward”. In: *Ecological Economics* 138, pp. 252–265.

Bezin, Emeline (2019). “The economics of green consumption, cultural transmission and sustainable technological change”. In: *Journal of Economic Theory* 181, pp. 497–546.

Boyd, Robert and Peter J Richerson (1988). *Culture and the evolutionary process*. University of Chicago press.

Buenstorf, Guido and Christian Cordes (2008). “Can sustainable consumption be learned? A model of cultural evolution”. In: *Ecological Economics* 67.4, pp. 646–657.

Castellano, Claudio, Matteo Marsili, and Alessandro Vespignani (2000). “Nonequilibrium phase transition in a model for social influence”. In: *Physical Review Letters* 85.16, p. 3536.

Castro, Juana, Stefan Drews, Filippos Exadaktylos, Joël Foramitti, Franziska Klein, Théo Konc, Ivan Savin, and Jeroen van den Bergh (2020). “A review of agent-based modeling of climate-energy policy”. In: *Wiley Interdisciplinary Reviews: Climate Change* 11.4, e647.

Davis, Taylor, Erin P Hennes, and Leigh Raymond (2018). “Cultural evolution of normative motivations for sustainable behaviour”. In: *Nature Sustainability* 1.5, pp. 218–224.

Deffuant, Guillaume, Frédéric Amblard, Gérard Weisbuch, and Thierry Faure (2002). “How can extremism prevail? A study based on the relative agreement interaction model”. In: *Journal of artificial societies and social simulation* 5.4.

Di Francesco, Tommaso and Daniel Torren-Peraire (2024). “(Mis) information diffusion and the financial market”. In: *arXiv preprint arXiv:2412.16269*.

Durham, William H (1991). *Coevolution: Genes, culture, and human diversity*. Stanford University Press.

Eppstein, Margaret J, David K Grover, Jeffrey S Marshall, and Donna M Rizzo (2011). “An agent-based model to study market penetration of plug-in hybrid electric vehicles”. In: *Energy Policy* 39.6, pp. 3789–3802.

Epstein, Joshua M (1998). “Zones of cooperation in demographic prisoner’s dilemma”. In: *Complexity* 4.2, pp. 36–48.

Epstein, Joshua M and Robert Axtell (1996). *Growing artificial societies: social science from the bottom up*. Brookings Institution Press.

Fan, Xia and Chuanju Li (2025). ““To be unfolding” or “be on its last legs”—Preferential tax policies between supply and demand and the development of the new energy vehicle industry based on an ABM model”. In: *Energy Policy* 200, p. 114552.

Glaeser, Edward L, Bruce I Sacerdote, and Jose A Scheinkman (2003). “The social multiplier”. In: *Journal of the European Economic Association* 1.2-3, pp. 345–353.

Greene, David L, Sangsoo Park, and Changzheng Liu (2014). “Analyzing the transition to electric drive vehicles in the US”. In: *Futures* 58, pp. 34–52.

Helfmann, Luzie, Jobst Heitzig, Péter Koltai, Jürgen Kurths, and Christof Schütte (2021). “Statistical analysis of tipping pathways in agent-based models”. In: *The European Physical Journal Special Topics* 230.16, pp. 3249–3271.

Herrmann, J.K. and Ivan Savin (2017). “Optimal Policy Identification: Insights from the German Electricity Market”. In: *Technological Forecasting & Social Change* 122, pp. 71–90.

Jaramillo, P., S. Kahn Ribeiro, P. Newman, S. Dhar, O. E. Diemuodeke, T. Kajino, D. S. Lee, S. B. Nugroho, X. Ou, A. Hammer Strømman, and J. Whitehead (2022). “Transport”. In: *Climate Change 2022: Mitigation of Climate Change*. Ed. by P. R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, and J. Malley. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK and New York, NY, USA: Cambridge University Press. Chap. 10.

Kaaronen, Roope Oskari and Nikita Strelkovskii (2020). “Cultural evolution of sustainable behaviors: Pro-environmental tipping points in an agent-based model”. In: *One Earth* 2.1, pp. 85–97.

Kauffman, Stuart A (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.

Konc, Théo, Ivan Savin, and Jeroen van den Bergh (2021). “The social multiplier of environmental policy: Application to carbon taxation”. In: *Journal of Environmental Economics and Management* 105, p. 102396.

Lacasse, Katherine (2016). "Don't be satisfied, identify! Strengthening positive spillover by connecting pro-environmental behaviors to an "environmentalist" label". In: *Journal of Environmental Psychology* 48, pp. 149–158.

Lamperti, Francesco, Giovanni Dosi, Mauro Napoletano, Andrea Roventini, and Alessandro Sapi (2018). "Faraway, so close: Coupled climate and economic dynamics in an agent-based integrated assessment model". In: *Ecological Economics* 150, pp. 315–339.

Mattauch, Linus, Cameron Hepburn, Fiona Spuler, and Nicholas Stern (2022). "The economics of climate change with endogenous preferences". In: *Resource and Energy Economics* 69, p. 101312.

McCoy, Daire and Seán Lyons (2014). "Consumer preferences and the influence of networks in electric vehicle diffusion: An agent-based microsimulation in Ireland". In: *Energy Research & Social Science* 3, pp. 89–101.

Mesoudi, Alex (2016). "Cultural evolution: a review of theory, findings and controversies". In: *Evolutionary biology* 43.4, pp. 481–497.

Niamir, Leila, Tatiana Filatova, Alexey Voinov, and Hans Bressers (2018). "Transition to low-carbon economy: Assessing cumulative impacts of individual behavioral changes". In: *Energy policy* 118, pp. 325–345.

Niamir, Leila, Gregor Kiesewetter, Fabian Wagner, Wolfgang Schöpp, Tatiana Filatova, Alexey Voinov, and Hans Bressers (2020). "Assessing the macroeconomic impacts of individual behavioral changes on carbon emissions". In: *Climatic change* 158.2, pp. 141–160.

Nuñez-Jimenez, Alejandro, Christof Knoeri, Joern Hoppmann, and Volker H. Hoffmann (2022). "Beyond Innovation and Deployment: Modeling the Impact of Technology-Push and Demand-Pull Policies in Germany's Solar Policy Mix". In: *Research Policy* 51, p. 104585.

Rai, Varun and Scott A Robinson (2015). "Agent-based modeling of energy technology adoption: Empirical integration of social, behavioral, economic, and environmental factors". In: *Environmental Modelling & Software* 70, pp. 163–177.

Read, Carveth (1914). *Logic, deductive and inductive*. A. Moring.

Rogers, Everett M. (2003). *Diffusion of Innovations*. 5th ed. New York: Free Press. ISBN: 978-0743222099.

Safarzyńska, Karolina and Jeroen van den Bergh (2022). "ABM-IAM: optimal climate policy under bounded rationality and multiple inequalities". In: *Environmental Research Letters* 17.9, p. 094022.

Savin, Ivan, Felix Creutzig, Tatiana Filatova, Joël Foramitti, Théo Konc, Leila Niamir, Karolina Safarzynska, and Jeroen van den Bergh (2023). "Agent-based modeling to integrate elements from different disciplines for ambitious climate policy". In: *Wiley Interdisciplinary Reviews: Climate Change* 14.2, e811.

Schumacher, Ingmar (2015). "The endogenous formation of an environmental culture". In: *European Economic Review* 76, pp. 200–221.

Silvia, Chris and Rachel M. Krause (2016). "Assessing the impact of policy interventions on the adoption of plug-in electric vehicles: An agent-based model". In: *Energy Policy* 96, pp. 105–118. ISSN: 0301-4215.

Wang, Tiantian, Bo Shen, Cecilia Han Springer, and Jing Hou (2021). "What prevents us from taking low-carbon actions? A comprehensive review of influencing factors affecting low-carbon behaviors". In: *Energy Research & Social Science* 71, p. 101844.

Windrum, Paul, Tommaso Ciarli, and Chris Birchenhall (2009). “Consumer heterogeneity and the development of environmentally friendly technologies”. In: *Technological Forecasting and Social Change* 76.4, pp. 533–551.

Wynes, Seth and Kimberly A Nicholas (2017). “The climate mitigation gap: education and government recommendations miss the most effective individual actions”. In: *Environmental Research Letters* 12.7, p. 074024.

Chapter 2

An agent-based model of cultural change for a low-carbon transition*

2.1 Introduction

The behavioural choices that compose an individual's lifestyle can greatly affect their carbon footprint (Wynes and Nicholas, 2017). Therefore, changes in individual consumption form an important part of reducing global greenhouse gas emissions. Theoretical models of behavioural change can provide insight into what barriers and drivers exist to the adoption of low-carbon lifestyles (Niamir, Kiesewetter, et al., 2020). A crucial element in the study of a transition towards low-carbon lifestyles is how quickly behaviours in a population, and even culture itself, will change. A good understanding of the relation between culture and behaviours can help to avoid potential lock-in towards brown alternatives (Buenstorf and Cordes, 2008; Burton and Farstad, 2020). Especially of interest is building an awareness of what cultural barriers may appear due to said interaction (Carattini et al., 2018).

Models of opinion dynamics can give insight into associated processes of polarization versus consensus formation (Deffuant, Amblard, et al., 2002), and transitions between these regimes (Helfmann et al., 2021; Castellano, Marsili, et al., 2000). The descriptive power of these models can be further enhanced through the inclusion of frameworks such as cultural evolution. This is particularly true for the role of repeated social learning amongst direct, oblique and peer connections in a social network and of biased transmission of behaviours. The first and second cases represent the influence of overlapping generations, whilst the second and third reflect the impact of non-direct-blood relations, such as with wider members of a community.

The application of cultural evolution to the issue of transition studies is a nascent area of research (Davis et al., 2018; Kaaronen and Strelkovskii, 2020; Buenstorf and Cordes, 2008). It lacks detail on the spread of multiple traits simultaneously over the same population. A multi-dimensional perspective is important due to the breadth of lifestyle changes required for deep decarbonisation. For example, Andersson and Nässén (2022) give empirical evidence for positive spillovers between intra-personal behaviours such as choosing to not fly, not owning a car, following a vegan diet, and not owning a semi-detached house. However not all environmentally related behaviours are equally

*This chapter was published as: Torren-Peraire, D., Savin, I., and van den Bergh, J. (2024). An agent-based model of cultural change for a low-carbon transition. *Journal of Artificial Societies and Social Simulation*, 27(1):13.

subject to social influence. Furthermore, this multi-dimensional approach centres the role of inter-behaviour spillovers in adoption dynamics, where the uptake of one green behaviour may lead to a self-perception of a greater green identity (Lacasse, 2016), with a stronger environmental identity making pro-environmental behaviour more probable (Van der Werff et al., 2014).

Our primary research question is: how does pro-environmental diffusion of behaviour interact with longer-term cultural evolution on a path to a low-carbon economy? We address this through answering three sub-questions: Firstly, how do culture and behavioural diffusion interact? Secondly, what is the longer-term impact of cultural change? Thirdly, what are the mechanisms behind behavioural decarbonisation? To study social interactions, which include bounded rationality of individuals, the use of agent-based models (ABM) is appropriate (Railsback and Grimm, 2019; Castro et al., 2020; Savin et al., 2022). Agents have internal properties as well as rules of interaction. These models reproduce complex behaviour from local microeconomic interactions (Epstein and Axtell, 1996), which allows for the simulation of a wide variety of phenomena (Kaaronen and Strelkovskii, 2020; Waring et al., 2015; Rai and Robinson, 2015; Kraan et al., 2019) including the diffusion of cultural traits in a population (Axelrod, 1997). Thus, ABMs lend themselves well to the study of cultural and behavioural change as these processes occur through the progressive accumulation of social interactions.

2.2 Cultural change and identity

Definitions of culture are almost as wide-ranging as the heterogeneity found in the culture itself. To produce an informative model of lifestyle change we require an instrumental and easily interpretable definition of culture. A core element of many definitions is social spreading, with Durham (1991) describing culture as a “system of symbolically encoded conceptual phenomena that are socially and historically transmitted within and between populations”. Similarly, evolutionary (Boyd and Richerson, 1988; Henrich and McElreath, 2003; Mesoudi, 2016), economic (Bezin, 2019; Bisin and Verdier, 2001) or physics-based models (Axelrod, 1997; Epstein and Axtell, 1996; Kuperman, 2006) represent culture as an abstract property or trait of an individual which can spread in a population. Whilst this definition may lack consideration of factors such as geographic location (Gupta and Ferguson, 1992) and how the micro-process of cultural transmission occurs (Kashima, 2008), it narrows the scope of what interactions, information or objects may be considered as culture.

Culture can provide the framework within which “strategies to respond to problems are devised and implemented” (Adger et al., 2013). This response component is especially of interest when considering solutions to the climate crisis. Information bubble filters (Geschke et al., 2019) or false consensus biases (Drews et al., 2022) can affect to whom, and what, individuals pay attention. This can slow down social tipping processes if there is a disconnect between the understood and real consequences of current pro-environmental behaviours (Wynes and Nicholas, 2017). Due to these biases, there may be heterogeneity in the quantity of information and length of exposure individuals require to change course concerning their environmental identity. This resistance to respond to new stimuli can be conceived of as cultural inertia. These preferences may change much more slowly over time thereby limiting the effectiveness of climate policy (Davis et al., 2018), achieving only shallow decarbonisation.

To study the impact of culture on an individual’s behaviour we require a model of how their culture changes over time. Cultural evolution is a Darwinian process explaining how long-term population-level changes in culture occur (Boyd and Richerson, 1988). It is constructed from three components: variation, fitness and inheritance. Cultural traits such as words and ideas exhibit variation (principle of variation). These in turn have differing rates of reproduction or transmission depending on the environment (principle of fitness). Finally, the traits present in a population provide a pool from which new generations can learn. As a result of social learning, there is a correlation between historical cultural traits and those of the next generation (principle of inheritance). Social learning refers to knowledge acquisition through the involvement of others, e.g. imitation (Mesoudi, 2016).

Existing computational models of cultural evolution can describe the spread of a single cultural trait (Boyd and Richerson, 1988; Henrich, 2001; Kaaronen and Strelkovskii, 2020), where an agent’s cultural parameter may change through individual and social learning, as well as through interactions with cultural niche infrastructure. Alternatively, the non-cultural-evolutionary model of Epstein and Axtell (1996) represents a “cultural chromosome” in the form of a set-length binary string called a cultural tag. This multi-dimensional tag represents a set of cultural attributes that an agent possesses. Here agents belong to one of two cultural groups depending on the majority of either 1s or 0s in the cultural tag. It is important to note that there is no consideration of where in the string these digits lie, only the quantity of each value matters. Thus two agents may belong to the same cultural group with a very distinct set of cultural attributes. Furthermore, Axelrod (1997) assigns culture as “the set of individual attributes that are subject to social influence”. Here, agents are limited with whom they can interact based on a vector of cultural features, each feature having a set of possible discrete values.

To produce a more specific and verifiable model of cultural change we further narrow our focus to studying the change of identities relevant to environmental behaviours. The background of an individual’s environmental identity can affect the decisions made regarding whether to engage or not, in certain behaviours (Van der Werff et al., 2013). Identities are independent of the behavioural outcomes of an individual’s decision-making process. Instead, they are self-defining, such that two agents may behave very differently but may identify themselves with the same group if they hold the same attitude towards said behaviours (Smaldino, 2019).

When considering several related behaviours, we take the approximation that identity is an outgrowth of culture (Grimson, 2010). This reduces the scope of what is necessary to consider when modelling the dynamics of culture and limits the generality of our model as it only applies in conditions of proximity between behaviours when one might expect a person to be acting under the same identity. Additionally, this approximation facilitates the comparison and use of empirical data regarding environmental identities and attitudes, such as in Nigbur et al. (2010). In contrast, validation and parameterisation of a purely cultural model would require more abstract data which would be harder to measure given the broader scope of the subject. Identities are not fixed and may change due to different contexts and evolution over time (Fielding and Hornsey, 2016). Additionally, individuals cannot express their full identities in social interactions as these are too rich; only certain subsets are represented, depending on the social context and interaction setting (Smaldino, 2019). We do not attempt to model how changing environments and decisions can cause particular identities to dominate in a behavioural decision process. Instead, we fix the context by only considering environmentally related behaviours, such

as deciding to reduce home energy use or whether to install domestic solar panels. This highlights the influence of several behaviours on an individual's environmental identity.

2.3 Model components

2.3.1 General structure of the model

To address our research questions, we produce an ABM of changes in individuals' lifestyles by considering their evolving behavioural choices. The model structure is shown in Figure 2.1. Individuals have a set of environmental behavioural traits that spread through a fixed Watts–Strogatz graph via social interactions with their ego-network. These exchanges are mediated by transmission biases informing from whom an individual learns and how much attention is paid. The influence of individuals on each other is a function of their similarity in environmental identity, where we represent environmental identity computationally by aggregating past agent attitudes towards multiple environmentally related behaviours. To perform a behaviour, agents must both have a sufficiently positive attitude towards a behaviour and overcome a corresponding threshold. This threshold structure, where the desire to perform a behaviour does not equal its enactment, allows for a lack of coherence between attitudes and actual emissions. This leads to a disconnect between what people believe and what they do, such that the social network as whole desires greener behaviours but only a minority performs them. Subsequently, we outline the justification and assumptions underlying model components.




Figure 2.1: *Model structure composed of behavioural components (yellow), opinion dynamics (blue), environmental identity (purple) and environmental output (green). Arrows indicate the direction of influence between components and stacked boxes represent multiple individuals (grey) and a behavioural vector (yellow and green).*

To evaluate the impact of individual decisions on overall emissions we require a measure of whether individuals act on their intentions and opinions towards “greener” behavioural alternatives. This process is represented in the yellow containers in Figure 2.1. An intention or favourable attitude towards a behavioural option is not sufficient; one must also have the means or control to do so according to the Theory of Planned Behaviour (TPB) (Ajzen, 1991). This theory may be used to study motivations behind

pro-environmental social behaviours, such as in Nigbur et al. (2010) which empirically investigates participation in recycling programmes. Alternatively, Niamir, Filatova, et al. (2018) use the TPB within an empirically based ABM, where agents must overcome several barriers along the road to undertaking one of three possible pro-environmental behaviours.

Taking inspiration from the TPB, we model the extent to or frequency with which a behaviour is performed as a balance between a socially influenced behavioural attitude against a static threshold. Each behaviour is represented as a one-dimensional continuous parameter between extremes of a zero-emissions green choice and maximally emissive brown reference. Individuals $n = \{1, \dots, N\}$, each have multiple environmentally relevant behaviours $m = \{1, \dots, M\}$ that evolve over discrete generic time t . These behaviours represent abstract environmentally related actions such as the decision to install solar energy panels, whether to choose a brown or green energy provider, cycling to work but only if weather conditions are suitable or flying less but not entirely stopping. For each behaviour m , the continuous value $B_{t,n,m} \in [-1, 1]$ is an index that represents the extent to which or frequency with which a behaviour is performed. Modelling this way is a compromise between representing the detail of context-dependent behavioural decisions and maintaining an abstract generality of the behaviours chosen. Put explicitly, the behavioural value $B_{t,n,m}$ is determined by two continuous variables: the individual's attitude $A_{t,n,m} \in [0, 1]$ towards the behaviour and the threshold or barrier of entry for performing a behaviour $T_{n,m} \in [0, 1]$. A value of $A_{t,n,m} = 1$ is the “greenest” attitude and $A_{t,n,m} = 0$ the “brownest” or most indifferent to environmental impacts, and similarly $T_{n,m} = 1$ is the highest barrier of entry and $T_{n,m} = 0$ the lowest. Therefore the behavioural value $B_{t,n,m}$ is given by

$$B_{t,n,m} = A_{t,n,m} - T_{n,m}, \quad (2.1)$$

where initial values $T_{n,m}$ and $A_{0,n,m}$ are generated separately using a Beta distribution, see Figure 2.A.1 in the Appendix. This was chosen due to the ease with which uniform, asymmetric and polarised distributions may be generated. The form of the Beta distribution is given by two parameters a and b . The expectation value is dictated by the ratio $a/(a + b)$, whilst the degree of polarisation is inversely proportional to the magnitude of a and b . We use a_A, b_A and a_T, b_T to describe the initial attitude and threshold distributions. Thus the larger the ratio of $a_A/(a_A + b_A)$ the “greener” initial attitudes. Conversely, the larger the ratio of $a_T/(a_T + b_T)$ the higher the threshold or barrier of entry for performing a behaviour. A value of $-1 \leq B_{t,n,m} \leq 0$ represents the brownest behavioural choice, whilst a value $0 < B_{t,n,m} \leq 1$ is a “greener” behaviour. To decrease model complexity, thresholds to performing behaviours, $T_{n,m}$ are heterogeneous between agents but static.

The total emissions E_t produced by the population of size N is given by the summation over each of the multiple behaviours performed by each individual,

$$E_t = \sum_{n=1}^N \sum_{m=1}^M \frac{1 - B_{t,n,m}}{2}, \quad (2.2)$$

where the form of the summand in Equation 2.2 ensures that a single perfectly green behavioural choice, $B_{t,n,m} = 1$, results in zero emissions for that m^{th} behaviour of the n^{th} individual. On the other hand, its brown counter part, $B_{t,n,m} = -1$, results in a single unit of emissions.

2.3.2 A dynamic model of culture using environmental identity

We represent three key aspects of identity. Firstly the slower, longer-term change process of environmental identity driven by a faster behavioural diffusion process, secondly, the central role of socially transmitted information and finally the cyclic, self-defining, nature of identity. Concerning this third point, agent identity is defined not through the behaviours an agent performs but instead through their opinion or attitude towards said behaviours. This self-defining process is highlighted in the lower loop of Figure 2.1. Here, we create bidirectional causal relations between agent attitudes and identity (Schaller and Muthukrishna, 2021) such that identity “becomes what the constituting agents make it to be” (Fáth and Sarvary, 2005). However, these attitudes towards environmentally related behaviours are themselves determined through social information exchanges, so that we conform to the definition previously laid out. Furthermore, the longer-term change process is captured through a weighted average of previous behavioural attitudes. How far this moving average reaches back in time is determined by a cultural inertia parameter, with a greater value meaning agents are influenced by past opinions for a longer time (Konc, Drews, et al., 2022). The weighting is given by a hyperbolic discounting factor (Loewenstein and Prelec, 1992; Laibson, 1997; Yi et al., 2006), such that current identities are influenced by recent history with diminishing importance the further back in time is considered.




Figure 2.2: *Discount array D_t as a function of time for different discount parameters δ .*

We adapt the cultural chromosome proposed by Epstein and Axtell (1996) considering multiple continuous attitudes towards related environmental behaviours that determine a single identity variable. We distinguish ourselves from single-parameter diffusion models by adding internal dynamics to our identity parameter through this mean behavioural attitude. Therefore identity is represented as an aggregate of all behavioural attitudes of an individual over time with a discount factor putting greater value on more recent behavioural attitudes. In particular we define the environmental identity I_n of agent n at time t as

$$I_{t,n} = \left[\frac{1}{\sum_{s=0}^{\rho} \delta^s} \right] \sum_{s=0}^{\rho} \delta^s \bar{A}_{t-s,n}, \quad (2.3)$$

where the mean behavioural attitude \bar{A} over M behaviours is given by

$$\bar{A}_{t,n} = \frac{1}{M} \sum_{m=1}^M A_{t,n,m}, \quad (2.4)$$

where a low attitude value contributes to an identity of indifference towards the environment. Additionally, ρ is a cultural inertia parameter representing the duration of the past considered, s is a dummy variable for the discrete-time step in the present or past evaluated, \bar{A} , the average attitude over one time step for M behaviours, $\delta \in [0, 1]$ is a discount parameter that produces a hyperbolic discounting effect. In the discount array $D_s = [1, \delta^1, \delta^2, \dots, \delta^\rho]$, each subsequent time step contributes a smaller amount to the weighted average. In Figure 2.2, the extent to which individuals are resistant to changing their identities is determined by the discounting array, D_t . The cultural inertia parameter $\rho = 100$. Note that the smaller the value of the discount parameter δ , the more aggressively the influence of past attitudes is discounted.

2.3.3 Information diffusion through imperfect imitation

Interactions within a social network play a greater role in changing attitudes towards more socially susceptible behaviours that contribute to conspicuous consumption. For example, the domestic installation of solar panels has been shown to increase the likelihood of adoption by others in neighbourhoods (Rai and Robinson, 2015; Bollinger and Gillingham, 2012; Baranzini et al., 2017; Carattini et al., 2018). On the other hand, behaviours that are less socially susceptible, such as reductions in domestic energy use, may be harder to influence through social interactions alone. However, even these behaviours may be nudged in a “greener” direction through the introduction of descriptive and injunctive social norms (Schultz et al., 2007; Davis et al., 2018). The attitude of individuals towards M behaviours, which each differ in social susceptibility ϕ_m , varies according to their past attitudes and the social influence of their ego-network. This means individuals form an attitude towards a given behaviour through a cumulative process of repeated social interactions. The evolution of individuals’ attitudes is given by

$$A_{t+1,n,m} = [1 - \phi_m] A_{t,n,m} + [\phi_m] S_{t,n,m}, \quad (2.5)$$

where the attitude to behaviour m is modulated by ϕ_m , a measure of conspicuous consumption, and $S_{t,n,m}$ is the social influence component due to an agent’s K_n ego-network members. In the extreme of $\phi_m = 0$ an agent’s attitude stops evolving. This represents a behaviour performed in isolation of social pressures. Conversely, if $\phi_m = 1$ agent attitudes’ towards that behaviour are entirely determined through social interactions.

In forming opinions the influence of our ego-network is paramount as often one cannot rely on one’s judgement solely. Opinion dynamics models aim to explain how the attitude of an individual within a group may evolve through social interactions where opinions are exchanged. In constructing an equation for social dynamics of behavioural attitudes, we consider literature from non-Bayesian opinion dynamics models using finite social networks. For a detailed literature review of social influence and opinion dynamics models see Mason et al. (2007), Castellano, Fortunato, et al. (2009), Acemoglu and Ozdaglar (2011) and Grabisch and Rusinowska (2020).

Now we formally elaborate on the blue components in Figure 2.1 to model the influence of social learning on individuals’ attitudes towards environmentally relevant behaviours.

A key distinction between models is whether they use continuous or binary opinion parameters. Moreover, the issue of whether a single opinion or a high-dimensional vector of opinions is considered may also be used to draw lines of division. For example, Schelling (1969) and Granovetter (1978) are pre-eminent cases of binary single opinion models that, whilst consisting of simple inter-agent rules of interaction, describe complex social phenomena. On the other hand, DeGroot (1974) uses a continuous opinion for each agent which is updated in discrete time steps using a mean of all other agents.

We implement a multi-dimensional continuous model inspired by DeGroot (1974) to represent social learning in our model, with a weighted mean to aggregate the impact of each agent's ego-network. To model the imperfection of social transmission we add a Gaussian error $\varepsilon = G(0, \sigma_\varepsilon^2)$. A separate transmission error is applied to each of the M behaviours after an individual has aggregated the attitudes of their ego-network. Thus, the social learning component is given by

$$S_{t,n,m} = \left[\sum_{k=1}^{K_n} \alpha_{t,n,k} A_{t,k,m} \right] + \varepsilon_{t,n,m}, \quad (2.6)$$

where K_n is the total number of agents in the n^{th} agent's ego-network. This is weighted by $\alpha_{n,k}$, which represents how much focal agent n values the opinion of agent k in their ego-network. Finally, $A_{t,m,k}$ is the positive attitude of agent k towards behaviour m . Our social learning error is pseudo-Gaussian, since we clip the value of $S_{t,n,m}$, such that $S_{t,n,m} \in [0, 1]$, this ensures that $A_{t,n,m} \in [0, 1]$.

2.3.4 Homophily and asymmetric weightings in a social network

In the model, social interactions occur between individuals within a Watts–Strogatz graph with a small-world property and a mean number of connections K per individual (Watts and Strogatz, 1998). This choice of graph is due to their pervasive nature in real-world physical social networks. Watts–Strogatz graphs begin by placing N individuals in a ring and then attaching K links between an individual and those other individuals closest to them. In this state the network has an extremely high degree of homophily and clustering as ego-networks overlap significantly. Subsequently, using a probabilistic re-wiring p_r of connections between two adjacent nodes to a third long-distance one we can introduce long-distance or weak ties (Granovetter, 1973). This greatly reduces the average shortest path length while only marginally reducing the clustering of the network, achieving the small-world property.

We generate attribute homophily in the network by placing individuals next to those with initial environmental identities that are most similar (Kapeller et al., 2019). The attribute homophily parameter h measures to what degree an individual's position in the network is a function of their identity. In the case of perfect attribute homophily, individuals are surrounded by ego-networks that are like-minded in terms of environmental identity. However, for heterogeneous initial identity values and Watts–Strogatz network, this forms a circle with an approximately constant gradient in identity, accounting for stochastic variations from the generating Beta distribution. This is shown in the third column of the second row of Figure 2.6. The degree of attribute homophily may then be varied discretely using a Fisher-Yates shuffle algorithm (Fisher and Yates, 1953), where random pairs of individuals swap places in the network. For our implementation, an attribute homophily parameter $h = 1$ means no pair swaps occur, whilst $h = 0$ results in N random swaps.

Boyd and Richerson (1988) model the notion that people's choices of whom to learn new traits from are not random. A confirmation bias is content-dependent where more attention is paid to information that reinforces current opinions (Lord et al., 1979; Nickerson, 1998). Furthermore, individuals that identify with a particular group tend to assimilate their norms. In a parallel fashion, they distance themselves from out-group individuals (Fielding and Hornsey, 2016). This forms a key component of the model by Axelrod (1997), where individuals interact more with those who are similar to them. This causes them to share more similarities and hence leads to further interactions. In the bounded confidence (Hegselmann, Krause, et al., 2002) or relative agreement (Deffuant, Amblard, et al., 2002) models, proximity in opinion or distance may be used to represent this phenomenon. Likewise, Konc and Savin (2019) use an opinion distance which is framed as a confirmation bias. Furthermore, in Deffuant, Neau, et al. (2000), agents may only interact if their vectors of opinions are sufficiently alike. They highlight the parallels with genetic transmission processes in that "reproduction only occurs if genome distance is smaller than a given threshold". Limitations in contact between groups can allow for minority opinions or ideologies to persist amongst an overwhelming majority (Bergh et al., 2019).

We replicate the stylised fact of interactions leading to further interactions found in Axelrod (1997) by basing the strength of agent opinion exchange on the similarity of their environmental identity. Following Konc and Savin (2019), we use a combination of an exponential form and confirmation bias to vary how much agents may interact with those of distinct identities. Greater values of confirmation bias mean that an agent pays attention to a rapidly shrinking sub-network within the ego-network. Partly inspired by Brock and Hommes (1998) use of discrete choice models, our social network weighting matrix $\alpha_{t,n,k}$ is given by the softmax function

$$\alpha_{t,n,k} = \frac{e^{-\theta|I_{t,n} - I_{t,k}|}}{\sum_{j \neq n}^{K_n} e^{-\theta|I_{t,n} - I_{t,j}|}}, \quad (2.7)$$

where we consider the identity distance of K_n agents within an ego-network and θ is a measure of confirmation bias. Thus, if I_n and I_k grow further apart over time, the influence they exert on each other also decreases¹. To ensure that the total weighting in an agent's ego-network is one we normalise the values against the total weighting calculated. Note that the social network links are static. Therefore even if there exist agents who are similar in identity, they will not be able to form new links. However, they may still influence each other through a shared ego-network. We choose this static network to ensure that the small-world properties used in the initial network are conserved. Therefore we also do not include the re-wiring probability as part of our sensitivity analysis.

¹The form of Equation 3.11 takes inspiration from a Boltzmann distribution $e^{-\beta E}/Z$ (McQuarrie, 2000), whereby the probability distribution of a statistical system occupying a state is proportional to the energy E , of that state, and is normalised by the partition function Z . In our model, the energy of a state is given by $|I_n - I_k|$. Moreover, thermodynamic β is inversely proportional to temperature. With increasing temperature, particles have greater kinetic energy, hence can move around more and "interact" with other particles. From this perspective, the confirmation bias θ can be seen as analogous to the inverse of temperature, such that increasing confirmation bias represents a system where individuals are "colder" to those different from them and pay less attention to their opinion. Whilst low values of confirmation bias correlate to an open society where individuals interact uniformly with their ego-network.

The external influence of an agent in an ego-network can change an individual's attitude towards a certain environmentally related behaviour. This in turn leads to a changing environmental identity which strengthens this social relation, leading to further behavioural change. The intermediary role played by identity can produce new behavioural dynamics compared to a reference case in which behavioural attitudes evolve through independent social interactions. The greater the degree to which a behaviour is socially influenced, determined by ϕ_m , the more it is susceptible to outside shocks, driving behavioural change.

2.3.5 Baseline experimental set-up

When exploring model dynamics we typically study the case of $N = 200$ individuals for $\tau = 3000$ time steps. Note that we envision a single period being between one week and a day, however, for a theoretical model like this, is not strictly necessary. A Watts–Strogatz network is used with a mean node degree $K = 20$ giving a typical network density ≈ 0.1 (the ratio of actual to potential inter-agent links). The parameters required to run the model, and the test case ranges explored in the sensitivity analysis, are summarised in Table 2.1 with variables shown in Table 2.2.

It is important to note that given the softmax function form of Equation 3.11 for inter-agent weighting there is always some influence from an individual's dis-similar ego-network. This means that given sufficient time (and low social learning error or perfect imitation) the population will reach a single consensus identity. However, the immediacy with which a transition to low-carbon lifestyles must occur does not allow for such an extended time frame. Therefore for all experimental runs, we fix the number of simulated timesteps at $\tau = 3000$, as we are specifically interested in the identity dynamics of the model in short periods and the path dependency of consensus formation in the network. Furthermore, the rate at which behavioural change occurs is dictated by the conspicuous consumption parameter ϕ_m , with lower values leading to slower change. We assign different ϕ_m values in the range $[0.01, 0.05]$ to each of the M behaviours, this represents the varying social susceptibility of environmentally relevant behaviours. Therefore we exclude both the number of time steps and the conspicuous consumption parameter ϕ_m from the sensitivity analysis.

Table 2.1: List of model parameters including ranges explored or tested in sensitivity analysis and other experiments.

Parameter Name	Symbol	Definition	Range	Comments
Number of individuals	N	Total number of individuals in the social network	[100,1000]	
Number of behaviours	M	Behaviours modelled per agent	[1,30]	
Mean node degree	K	Mean number of members of ego-network	[5,99]	This must be less than the total possible number of connections given by $N - 1$.
Cultural inertia	ρ	Number of timesteps over which past attitude states influence current identity state	[1,3000]	A lower bound of 1 means that only the present is considered and an upper bound of 3000 is chosen as this must be $\leq \tau$.
Social learning error standard deviation	σ_ε	Standard deviation of Gaussian learning error representing the degree of perfect imitation	[0,1]	In proportion to the scale of the attitudes and thresholds of [0, 1].
Discount factor	δ	Decrease in relative importance between two adjacent moments in time	[0,1]	A value of 0 means only present attitudes affect identity, whilst 1 means all past timesteps considered equally.
Attribute homophily	h	Degree of identity homogeneity in the initial social network	[0,1]	$h = 1$ no pair swaps occur, $h = 0$ results in N random swaps.
Confirmation bias	θ	How much agents only listen to members of their ego-network with similar identities	[0,200]	Set the lower bound as 0 as negative values mean that individuals seek out identities opposite to their own.
Initial attitude and threshold Beta (a,b)	a_A, b_A, a_T, b_T	a and b are the two inputs for the Beta distribution to generate initial distribution of agent attitudes and thresholds	[0.05, 8]	Ranges allow for distributions representing polarisation and consensus in initial attitudes and thresholds.
Total time steps	τ	Discrete simulation calculations	3000	Not included in sensitivity analysis.
Conspicuous consumption factor	ϕ_m	Behaviour specific social susceptibility, determines simulation speed	0.01 - 0.05	Not included in sensitivity analysis to keep simulation speed constant.
Probability of re-wiring	p_r	Likelihood that connection between agents are swapped to form long-distance or weak tie	0.1	Not included in sensitivity analysis to conserve small-world property.

Table 2.2: List of model variables

Variable Name	Symbol	Definition
Attitude	$A_{t,n,m}$	How much an agent wants to perform behaviour m
Threshold	$T_{n,m}$	How high the barrier of entry is to performing behaviour m
Behavioural value	$B_{t,n,m}$	To what extent or with what frequency behaviour m is performed
Total emissions flow	E_t	Sum of emissions flow due to N individuals each with M behaviours
Identity	$I_{t,n,m}$	Degree to which an agent associates themselves with a pro-environmental identity
Social learning component	$S_{t,n,m}$	The influence of agent n 's ego-network on their attitude towards behaviour m
Social network weighting	$\alpha_{t,n,k}$	Matrix of inter-agent opinion importance
Node degree	K_n	Number of members of an individuals ego network. On average this is K but may vary due to network re-wiring

2.4 Results

We first give an overview of the model outcomes through a study of the typical identity dynamics towards three metastable states. Due to the model's complexity, exploring the entire parameter space is not feasible. Instead, we collate different kinds of phenomena produced by the model and explain them. In pursuit of this, we consider the processes of bifurcation, polarisation and consensus formation (Deffuant, Amblard, et al., 2002). Therefore, we also take a more detailed look at what model dynamics are induced by specific parameters. Additionally, we consider the effect of green influencers that act as fountains of green attitudes in the model, focusing on which components dictate the degree of behavioural decarbonisation. Finally, this is complemented by a sensitivity analysis to identify which parameters have the most impact on key outcome variables such as the variance in final identities of individuals, total emissions and relative change in emissions between the start and finish of an experiment.

2.4.1 Environmental identity dynamics

The identity dynamics produced by the model may be divided into three states as a function of the variance of final identities in the population. Examples of these different dynamics are shown in Figure 2.3, where each case is a time series of the identity dynamics of $N = 200$ individuals. The experiments differ through variations of the Beta a_A and b_A parameters for individuals' initial behavioural attitudes and confirmation bias θ . Case A ($a_A = 2.0$, $b_A = 2.0$, $\theta = 10$) represents the simplest model outcome, where approximatively normally distributed attribute values produced by large Beta distribution parameters, $a_A, b_A > 1$, lead to rapid consensus formation around a single population environmental identity. Decreasing Beta parameter values produce greater polarisation in initial conditions, in combination with large values of confirmation bias $\theta > 10$, this leads to the formation of splinter identity subgroups within the population. These can form a two-identity state metastable state as can be seen in case B ($a_A = 0.3$, $b_A = 0.3$,

$\theta = 18$). The greater the distance between these two-identity subgroups the slower the process of reconciliation occurs, and the greater the time frame required for consensus formation. Finally for sufficiently low values $a_A, b_A < 0.1$ and large θ , see case C ($a_A = 0.05$, $b_A = 0.05$, $\theta = 40$), the population remains splintered in multiple information or “identity bubbles” of individuals who only interact with a small group. A more detailed breakdown of the relation between a_A, b_A and the total emissions may be found in the Appendix Figure 2.A.2, where lines of constant expectation value of the Beta distribution $a_A/(a_A + b_A)$ are proportional to a constant level of emissions.

To further explore the effect of confirmation bias on identity dynamics we look at the bifurcation process of clusters of behavioural attitudes at the end of experiments. We use a Gaussian kernel density estimator to group individuals, measuring the location of these for increasing confirmation bias θ . The transitions between the three cases identified in Figure 2.3, are modulated by θ , as shown in the left of Figure 2.4. In the sub-figure we consider the effect of increasing confirmation bias on the location of final attitude clusters of the first, $m = 1$, of a total of three behaviours, $M = 3$. The location of these attitude clusters is determined using a Gaussian kernel density estimator with a bandwidth of 0.01. All experiments use the same initial seed to account for stochastic effects. For the same degree of polarisation in initial attitudes, larger values of θ produce greater identity fragmentation. Sufficiently high values of θ lead to splintering, in the style of case C.




Figure 2.3: Time series of individual’s environmental identity, for a population of $N = 200$, where the greater $I_{t,n}$, the “greener” a person’s identity. Increasing (A to C) initial attitude polarisation and confirmation bias creates strands of “identity bubbles” and slower consensus formation.

A key cornerstone of the model is the role of cultural identity in the diffusion of pro-environmental behaviours. To investigate this effect we perform the same bifurcation analysis in the case of behavioural independence, see right Figure 2.4. For these experiments, the social network weighting $\alpha_{t,n,k}$ is now determined by the behavioural attitude distance, $|A_{t,n,m} - A_{t,k,m}|$, not the identity as in Equation 3.11. This results in one weighting matrix for each M behaviours, $\alpha_{t,n,k,m}$. In this scenario, the fragmentation

in attitudes occurs at much lower values of confirmation bias. Moreover, for these low confirmation bias values, identity allows for the formation of larger, more behaviourally heterogeneous, groups relative to the behavioural independence case. Therefore identity stimulates the convergence of opinions by allowing individuals of more diverse behavioural backgrounds to relate themselves better to their peers and imitate their behaviour.




Figure 2.4: *Bifurcation diagram showing the effect of increasing confirmation bias on the location of final attitude clusters of the first, $m = 1$, of a total of three behaviours, $M = 3$, for behavioural inter-dependence (environmental identity) and independence cases. Environmental identity stimulates the convergence of attitudes, relative to the behavioural independence case.*

2.4.2 Impact of model components

To study how the frequency with which individuals update their identity can impact identity dynamics, we consider three cases. The first case, static uniform weighting, represents a society in which individuals value the opinion of all their ego-network members in the Watts–Strogatz network equally such that $\alpha_{t,n,k} = 1/K_n$ and are unable to change their weighting over time. This is equivalent to having no confirmation bias, $\theta = 0$. In the second case, static culturally determined weighting, individuals calculate their social network weighting once based on their initial identities, according to Equation 3.11. Similarly to the first case, this is fixed for subsequent time periods of the experiment. In the third case, dynamic cultural weighting, we update $\alpha_{t,n,k}$ every time step, representing frequent social interactions. The columns in Figure 2.5 correspond to these three scenarios sequentially, whilst the top row gives the identity time series and the bottom row the step social network weighting matrix at the final time step. Each experiment is run for $N = 50$ individuals so that the heterogeneity in the social network weighting matrix is more visible.




Figure 2.5: Individuals' identity time series and final weighting matrices for three scenarios of differing identity updating frequency. Environmental identity consensus formation occurs slower when individuals can update social interaction weightings frequently as they can form "identity bubbles".

For the first and second cases, individuals reach a population identity consensus faster due to cross networks connections that "infect" each other with opposing views. In both cases individuals cannot update their social network weighting, meaning that new socially acquired information (which may conflict with their current attitudes) cannot be ignored. This demonstrates the strength of weak ties in breaking homophily effects. However, the pace of consensus formation in the second case is much slower as individuals pay less attention to weak ties, thus slowing the behavioural diffusion process. This is reflected in the $\alpha_{n,k}$ values for the second column and row of Figure 2.5 where agents have very specific individuals to whom they pay attention. For frequent updating of $\alpha_{t,n,k}$ agents can form "identity bubbles" which leads to the grouping of agents into strands that block out information resulting in a fractured identity spectrum by the end of the simulation. Interestingly, the social network weighting matrix at the final time step for this third case is similar to that of the first, uniform, weighting case. This suggests that these "identity bubbles" are within themselves uniformly distributed, with individuals paying equal attention to their peers but excluding those from the out-group, see Figure 2.A.3 in the Appendix. It is this stark distinction between in- and out-group individual weightings that sustains the identity strands and prevents global-, whilst enforcing local-, consensus formation.

Decreasing initial identity similarity between ego-network members leads to faster consensus formation as agents are exposed to those with distinct views, see Figure 2.6. Three experiments are run for identical initial conditions, crucially including graph structure, varying solely on the attribute homophily parameter h which dictates how mixed ego-networks are in the initial social network. A population size of $N = 100$ is used to highlight the differences in the initial identity network layout. Furthermore, the network structure is also the same, thus it is purely through reduced homophily that a societal identity consensus is reached faster, and not a greater presence of weak ties. This effect is stronger for greater initial values of attitude polarisation with $a_A, b_A < 0.5$, as the initial identity distance between the "greener" and "browner" identity groupings is greater. Therefore by facilitating the mobility of individuals and exchange between people with distinct environmental views, to break up social network homophily, policy-makers can

further foster consensus formation in pro-environmental behavioural choices.




Figure 2.6: *Decreasing initial identity homophily results in faster consensus formation in the network. Smaller values of attribute homophily h result in a more mixed ego-network, whilst $h = 1$ means individuals are placed next to those who have the most similar identities to them.*

Greater values of the discount factor δ mean that there is a focus on recent events, thus changes in the model can occur faster as there is less influence from past behavioural attitudes. Variation in the cultural inertia parameter ρ was found to have little effect on the model culture. This lack of impact is because additional time steps included in the discount array, through greater values of ρ , have an exponentially decreasing effect. Hence the cumulative sum of their influence can be negligible. Moreover, due to the autoregressive nature of Equation 2.5 current values of attitudes are strongly correlated with previous time steps. Therefore, larger values of ρ introduce relatively little extra variance into the past attitudes vector averaged in Equation 2.3.

2.4.3 Attitude change through green influencers

Up to now, we have explored identity dynamics under a purely diffusive regime. In the following let us consider how the model behaves when we introduce green influencers. These are modelled as a minority of individuals who actively promote green lifestyles (Chwialkowska, 2019). We represent green influencers as having one behaviour, out of three ($m = 1, M = 3$), which is not susceptible to social influence with a perfectly green attitude $A_{t,n,1} = 1$, but behave as non-influencers individuals otherwise. Their inclusion increases the population by 10%, from 200 to 220. To account for the larger population size we proportionally increase the average number of mean ego-network members, from $K = 20$ to 22, to maintain a constant network density.




Figure 2.7: *Total emissions of non-influencers individuals (top panel) in societal scenarios with and without both green influencers and inter-behavioural dependency (environmental identity), for different societal confirmation biases (columns). The bottom panel shows the relative percentage change in emissions between scenarios with and without green influencers, for cases with (black) and without (red) inter-behavioural dependency.*

To assess the impact of these green influencers on final total emissions, and to exclude the drop in emissions due to simply introducing “greener” behavioural attitudes, we only measure behavioural emissions from non-influencers individuals. In the top panel of Figure 2.7, we consider four scenarios, varying whether or not to include green influencers and whether to include environmental identity (black) or not (red), through inter-behavioural dependence. The changes in inter-behavioural dependence are performed similarly to the experiments in Figure 2.4. Along the horizontal axis, we vary the mean initial attitudes of the population from “greener” to “browner”. In the case of including green influencers, this represents an increasing distance between the average attitude of non-influencers and that of perfectly green influencers. These scenarios are then run for low ($\theta = 5$) and high ($\theta = 20$) confirmation bias shown in the left and right columns. Moreover, in the bottom panel of Figure 2.7, we measure the relative change in emissions between experiments with and without green influencers, comparing the impact of including environmental identity (inter-behavioural dependence, black) or excluding it (behavioural independence, red). Note that we measure the intra-stochastic value emissions change. In both the top and

bottom panels of Figure 2.7, the shaded area gives the maximum and minimum values of the measured attribute for 10 different initial stochastic values.

Beginning with the top panel of Figure 2.7, we can see that total emissions of non-influencers are primarily determined by their initial attitude preferences, as identified in the Appendix Figure 2.A.2. The bottom panel of Figure 2.7 indicates that the emissions reduction achieved by the introduction of green influencers is a function of the mean attitude distance, confirmation bias and behavioural interdependency. Moreover, this panel demonstrates the difficulty in achieving deep behavioural decarbonisation. The degree of emissions reduction decreases sharply as the attitude distance between green individuals and the rest of the population approaches zero. Therefore, deep decarbonisation was not achieved through the presence of green influencers alone due to the heterogeneous static behavioural thresholds of individuals.

The greater the distance between green influencers and the mean initial attitude of non-influencers individuals the larger the potential for behavioural emissions reduction. Furthermore, larger initial attitude distances lead to lower inter-agent weightings of green influencers, as defined in Equation 3.11. The combination of these two countervailing forces, greater decarbonisation potential but the lower weighting of green influencers in social interactions, at larger attitude distances between green influencers and non-influencers individuals, leads to a U-shaped curve in emissions reduction. In the case of inter-behavioural dependency, the peak of the curve occurs at an attitude distance of 0.4 and 0.15 for the low and high confirmation bias cases respectively.

In contrast, when there is behavioural independence, green influencers are too distant in the attitude space to have significant inter-agent weights with non-influencers individuals. In the case of behavioural inter-dependence, non-influencers are willing to listen to green influencers due to similarities in the non-green behaviours ($m = 2, 3$). Green influencers can exploit this similarity in the environmental identity aggregate to spread their message. Additionally, under lower confirmation biases individuals are willing to listen to ego-network members with greater identity or attitude distance. With both behavioural inter-dependence and low confirmation biases this results in green influencers inducing decarbonisation over a broader range of attitude distances, as shown in the bottom panel of Figure 2.7. Moreover, because of this wide-ranging influence, the degree of decarbonisation is also stronger, as individuals feel the “pull” of green influencers further away in the attitude space. This results in a greater total number of impactful social interactions between them over the simulation period.

The limited behavioural decarbonisation achieved by the minority of green influencers (at best 5 – 12% emissions reduction) indicates the limit to what voluntary actions can achieve. This highlights the need for climate policies, such as carbon pricing or industry standards, which would reduce thresholds or barriers of entry to performing “greener” behaviours. Therefore, further research is required into the impact of said policies on behavioural emissions. Specifically, in the case of market-based instruments with incomplete emissions coverage, such as the EU-ETS (Foramitti et al., 2021), which may heterogeneously affect multiple green behaviours. Especially of interest is how policies targeting behavioural thresholds might synergise with those spreading “greener” attitudes, either through green influencers or information provision policies.

The peak of emission reduction at low distance in environmental attitudes indicates the need for an individual-specific tailored approach when providing green information. This would avoid alienating individuals who might not react to information provision policies if they are too green. Instead, the messaging would adjust for an initially brown

but increasingly green society, thereby taking advantage of the emissions reduction peak. Jointly, this means that with greater behavioural inter-dependence and lower confirmation bias, an information provision policy would be targeting, and impacting, a greater part of society.

2.4.4 Sensitivity analysis

Sensitivity analysis reveals how much variation in specific inputs can affect output variation (Hamis et al., 2021), ensuring that results and conclusions drawn are contextualised (Ligmann-Zielinska et al., 2014). For this analysis, we use the Sobol (Sobol, 2001) and Saltelli (Saltelli, 2002) methods, implemented using the SALib python library (Herman and Usher, 2017), run over 76800 experiments. The key indicators we use are final total emissions E/NM , the variance in final identities of individuals σ_I^2 and the change in total emissions between the start and end of each experiment, $\Delta E/NM$. We normalise the total emissions E over agent number and behaviours to account for the scale effect on results of having a larger population that is more active. To account for the impact of stochasticity in the model (caused by initial attitude distribution, attribute homophily, re-wiring probability and imperfect social learning) we average output variables over multiple runs with different initial stochastic seeds.




Figure 2.8: *Sobol sensitivity analysis of model parameters, showing first order index for three key outputs: Final normalised total emissions, E/NM , the variance in final identities of the population σ_I^2 and final normalised changes in emissions, $\Delta E/NM$.*

Figure 2.8 shows the results of this sensitivity analysis, where a greater first Sobol index value represents the greater relative importance of one parameter over others. The most influential parameters in determining the final model emissions are the initial distribution of behavioural attitudes a_A, b_B and thresholds parameters a_T, b_T . This result is due to the linearity of Equations 2.1 and 2.2 in determining behavioural emission, and because our diffusive mechanism of social learning only allows for imitation and not innovation to more green (or brown) behaviours. The exact value of the final total

emissions is strongly correlated with the initial identity distribution given in the Appendix Figure 2.A.2.

The variance of final identities is dominated by the number of behaviours M , the social learning error ε and the confirmation bias θ . The fractional contribution of each behaviour to an individual's identity decreases with an increasing number of behaviours M . This means that for large M values, extreme attitudes have little effect on the identity of an individual. Hence, individuals can be similar in identity even if they differ substantially in a few specific behavioural attitudes, as their extremist attitudes are counter-balanced by their more moderate views in Equation 2.4. Alternatively put, for a fixed total time a larger number of behaviours M leads to faster consensus formation in identity, as individuals are influenced by attitudes from a wider range of people, since they cannot find close matches, so it is harder to form “identity bubbles”. In the case of social learning error, this simply increases the fundamental uncertainty each individual holds regarding their attitudes, which on the aggregate identity level leads to greater variability per individual, and by extension also at a population level. Greater confirmation bias θ also acts to slow the diffusion of opinions in the network as individuals tend to listen to those who are very close to them in identity space.

The change in total emissions over the simulation period is driven by the social learning error, as this is the only means through which new information may be introduced into the population. Consequently, through Equation 2.6, the greater N and M the more of this misinformation is introduced per time step. The role of social learning error in determining emissions changes indicates room for policy intervention in the form of information provision through advertising or education, which is left for further research.

We run a similar sensitivity analysis for the case of including green influencers with a range tested of [1,100] with 1920 experiments shown in the Appendix Figure 2.A.4. We find that a similar set of parameters dominate the first order index of total emissions E/NM and identity variance σ_I^2 as in Figure 2.8. However, in the case of emissions change $\Delta E/NM$ there is a lower sensitivity to the social learning error ε and greater importance of the number of behaviours M . With more behaviours, an individual's environmental identity becomes less connected to a single perfectly pro-environmental norm promoted by green influencers, making normal individuals more receptive to their pro-environmental attitude message. Note that in this sensitivity analysis both the density of the network and the initial preferences vary unlike in Figure 2.7.

2.5 Limitations and outlook

The representation of behaviours as a continuous parameter allows for a scale of emissions corresponding to the degree to which a pro-environmental behaviour is performed. However, behaviours such as investing in domestic solar energy are binary choices that occur once. Furthermore, the ability to perform one environmentally related behaviour can be dependent on the past completion of another. For example, the choice to cycle to work requires that an individual already owns a bike. Thus, future empirically-based applications of the general model may want to modify the definition of the behavioural value $B_{t,n,m}$ according to the nature of that behavioural choice.

Our model of environmental identity, as a time-weighted average, allows for the substitutability of behaviours. Consider the case of two individuals, each with two behaviours. The first has behavioural attitude values $A_{t,1,m} = [0.5, 0.5]$, and the second $A_{t,2,m} = [0, 1]$.

Given the current formulation of Equation 2.3, these two would have the same environmental identity. The fact that two individuals, one of whom is entirely apathetic towards environmentally related behaviors while the other is a green and brown extremist, can be grouped under the same identity raises the need for further model refinement. To address this issue, one potential solution is to employ k-nearest neighbor classification based on behavioral attitude vectors, grouping individuals and assigning a representative environmental identity value to each group. Alternatively, a mechanism of internal dissonance (Dalege et al., 2018) could avoid the repeated interaction of individuals with conflicting attitudes.

An aspect of model dynamics that falls outside of the scope of this study is the role of what type of information is socially learnt. In the case of the first, we chose the diffusion of environmentally related attitudes $A_{t,n,m}$ instead of the behavioural action $B_{t,n,m}$. The imitation of a behavioural action instead of attitude would result in which descriptive norms are the dominant mediator in social interactions.

In this study, we consider solely networks with a small-world property. The dynamic nature of the social network weighting between individuals then leads to the formation of "identity bubbles". Further study could investigate how different degree distributions in the social network can affect the propensity for these "identity bubbles" to occur. In addition, the impact on total emissions of the placement of "green influencers" within networks with high asymmetry in node degree, such as scale-free, could be an interesting research avenue.

We envision the model being used in conjunction with climate policies such as carbon pricing, industry standards or information provision. These would build on our study of green influencers, to produce dynamic behavioural thresholds and drive attitudes in the population towards a green consensus. Our model could inform how said policies leverage network effects or are inhibited by the "identity bubbles" previously highlighted. For example, the study of green influencers could be extended to include the decaying effect of information provision policies (Allcott and Rogers, 2014).

Additionally, moral licensing effects could be included, whereby one green behavioural choice can cause negative spillovers in subsequent environmentally related decisions. This is especially of interest given the role of inter-behavioural effects on network attitude consensus highlighted in the model. The emissions of individuals due to behavioural choices do not directly contribute to model dynamics, see Figure 2.1. Therefore the underlying mechanics of behavioural interdependency due to culture with imperfect social learning could be generalised to study other systems such as healthy lifestyles.

Future work might analyse case studies for a particular country, period or group. Especially of interest would be studies involving larger behavioural changes, such as the mass adoption of cycling in Copenhagen in a similar fashion to Kaaronen and Strelkovskii (2020), but with a wider focus on how the adoption of these low-carbon behaviours affected other environmentally related behaviours (Andersson and Nässén, 2022). This approach would provide empirical validation for our theoretical mode, complement the robustness of our conclusions derived from our large-scale sensitivity analysis of model parameters.

The introduction of exogenous dynamic thresholds might better represent the degree to which certain behavioural choices are cyclic, such as reduced domestic energy savings in winter or the increased inconvenience of cycling in harsh weather. Considering longer time scales, further research could introduce a missing link between behavioural emissions, environmental conditions and agent choices. This would make external conditions

endogenous to agents' choices and thus create a cycle of choices between actions and changes in the environment. This might lead to innovation towards "greener" behaviours and thus the emergence of pro-environmental culture in the model. This could lead to integration with the nascent literature on ABM approaches to Integrated Assessment Models (IAMs) (Lamperti et al., 2018; Safarzyńska and Bergh, 2022), motivated by a need to provide greater detail in representations of agent heterogeneity (De Cian et al., 2020) in the context of demand-side emissions reduction.

2.6 Conclusions

In this article, we have developed and analysed an agent-based model of cultural dynamics. It describes how the diffusion of attitudes towards environmentally related behaviours can interact with longer-term cultural change. To this end, the model incorporates a cultural evolutionary framework, where culture is defined as socially transmitted information. This is represented in the model as an environmental identity which consists of slow long-term change driven by a faster behavioural diffusion process, where we aggregate multiple behavioural attitudes of an individual over time with a discount factor. Individuals interact in a small-world Watts–Strogatz network through imperfect imitation of behavioural attitudes. The impact of culture is to mediate the strength of social interactions according to environmental identity similarity, inducing behavioural interdependence.

As our primary research question, we wanted to answer how pro-environmental diffusion of behaviour interacts with longer-term cultural evolution on a path to a low-carbon economy. How do culture and behavioural diffusion interact? What is the longer-term impact of cultural change? What are the mechanisms behind behavioural decarbonisation? Firstly, considering the interaction of culture and behavioural diffusion, we found that the presence of culture, as an environmental identity, helps stimulate consensus formation in behavioural attitudes in large groups, relative to the case of behavioural independence. This inter-behavioural dependence facilitates interactions between individuals who may differ in specific attitudes but hold similar identities in the aggregate of multiple behaviours. Secondly, the longer-term component of culture plays a small role in the model dynamics due to the autoregressive nature of how attitudes change over time. This meant that values of environmental identity in the more distant past were very similar to those in the near past, resulting in little impact of extending the period over which current environmental identity is evaluated. Thirdly, we find the extent of behavioural decarbonisation of individuals in the social network to be strongly dependent on the initial distribution of preferences in behavioural attitudes and thresholds. Imperfect social learning drives the change in total emissions between the start and end of experiments as it acts as the sole source of new information in the model. Our results indicate that the speed of consensus formation in environmental identity is strongly influenced by exposure to information from individuals with contrasting opinions. This may be derived from sources such as inter-behavioural spillovers, confirmation biases in social interactions or breaking of homophily effects.

We drive individuals' attitudes towards greener outcomes through the addition of influencers who in one behaviour act as broadcasters of a perfectly green attitude. The inclusion of culture led to greater decarbonisation, compared to the behavioural independence case. In this scenario, green influencers overcame large distances in attitude

between them and non-influencers, by leveraging similarities in the attitude aggregate to spread their message to a wider audience. The impact of green influencers was found to be greatest when the initial attitude distance between non-influencers and green influencers was small enough to allow the them to remain relevant in social interactions, but great enough for there to still be large behavioural decarbonisation potential through solely attitude change. This indicates the need for individual-specific information provision policies to avoid alienating those who might be inert to pro-environmental information if it is too green. Moreover, the failure of this green influencer minority in achieving deep decarbonisation, through solely voluntary action, indicates the need for further modelling to assess the impact of culture on carbon pricing or industry standards policies that would reduce barriers to performing “greener” behaviours.

Model Documentation

The model is implemented in Python 3.9, and is available through the CoMSES Computational Model Library as “An agent-based model of cultural change for a low-carbon transition”: <https://www.comses.net/codebases/88c233af-a3dd-4ba0-bb01-a3f5990e5ff0/releases/1.1.0/>.

References

Acemoglu, Daron and Asuman Ozdaglar (2011). “Opinion dynamics and learning in social networks”. In: *Dynamic Games and Applications* 1.1, pp. 3–49.

Adger, W Neil, Jon Barnett, Katrina Brown, Nadine Marshall, and Karen O’brien (2013). “Cultural dimensions of climate change impacts and adaptation”. In: *Nature climate change* 3.2, pp. 112–117.

Ajzen, Icek (1991). “The theory of planned behavior”. In: *Organizational behavior and human decision processes* 50.2, pp. 179–211.

Allcott, Hunt and Todd Rogers (2014). “The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation”. In: *American Economic Review* 104.10, pp. 3003–37.

Andersson, David and Jonas Nässén (2022). “Measuring the direct and indirect effects of low-carbon lifestyles using consumption data”. In: *Journal of Cleaner Production*, p. 135739.

Axelrod, Robert (1997). “The dissemination of culture: A model with local convergence and global polarization”. In: *Journal of conflict resolution* 41.2, pp. 203–226.

Baranzini, Andrea, Stefano Carattini, and Martin Peclat (July 2017). *What drives social contagion in the adoption of solar photovoltaic technology*. GRI Working Papers 270. Grantham Research Institute on Climate Change and the Environment.

Bergh, Jeroen CJM van den, Ivan Savin, and Stefan Drews (2019). “Evolution of opinions in the growth-vs-environment debate: Extended replicator dynamics”. In: *Futures* 109, pp. 84–100.

Bezin, Emeline (2019). “The economics of green consumption, cultural transmission and sustainable technological change”. In: *Journal of Economic Theory* 181, pp. 497–546.

Bisin, Alberto and Thierry Verdier (2001). “The economics of cultural transmission and the dynamics of preferences”. In: *Journal of Economic theory* 97.2, pp. 298–319.

Bollinger, Bryan and Kenneth Gillingham (2012). "Peer effects in the diffusion of solar photovoltaic panels". In: *Marketing Science* 31.6, pp. 900–912.

Boyd, Robert and Peter J Richerson (1988). *Culture and the evolutionary process*. University of Chicago press.

Brock, William A and Cars H Hommes (1998). "Heterogeneous beliefs and routes to chaos in a simple asset pricing model". In: *Journal of Economic dynamics and Control* 22.8-9, pp. 1235–1274.

Buenstorf, Guido and Christian Cordes (2008). "Can sustainable consumption be learned? A model of cultural evolution". In: *Ecological Economics* 67.4, pp. 646–657.

Burton, Rob JF and Maja Farstad (2020). "Cultural lock-in and mitigating greenhouse gas emissions: The case of dairy/beef farmers in Norway". In: *Sociologia Ruralis* 60.1, pp. 20–39.

Carattini, Stefano, Martin Péclat, and Andrea Baranzini (2018). *Social interactions and the adoption of solar PV: evidence from cultural borders*. Grantham Research Institute on Climate Change and the Environment.

Castellano, Claudio, Santo Fortunato, and Vittorio Loreto (2009). "Statistical physics of social dynamics". In: *Reviews of modern physics* 81.2, p. 591.

Castellano, Claudio, Matteo Marsili, and Alessandro Vespignani (2000). "Nonequilibrium phase transition in a model for social influence". In: *Physical Review Letters* 85.16, p. 3536.

Castro, Juana, Stefan Drews, Filippos Exadaktylos, Joël Foramitti, Franziska Klein, Théo Konc, Ivan Savin, and Jeroen van den Bergh (2020). "A review of agent-based modeling of climate-energy policy". In: *Wiley Interdisciplinary Reviews: Climate Change* 11.4, e647.

Chwialkowska, Agnieszka (2019). "How sustainability influencers drive green lifestyle adoption on social media: the process of green lifestyle adoption explained through the lenses of the minority influence model and social learning theory". In: *Management of Sustainable Development* 11.1, pp. 33–42.

Dalege, Jonas, Denny Borsboom, Frenk van Harreveld, and Han LJ van der Maas (2018). "The attitudinal entropy (AE) framework as a general theory of individual attitudes". In: *Psychological Inquiry* 29.4, pp. 175–193.

Davis, Taylor, Erin P Hennes, and Leigh Raymond (2018). "Cultural evolution of normative motivations for sustainable behaviour". In: *Nature Sustainability* 1.5, pp. 218–224.

De Cian, Enrica, Shouro Dasgupta, Andries F Hof, Mariësse AE van Sluisveld, Jonathan Köhler, Benjamin Pfluger, and Detlef P van Vuuren (2020). "Actors, decision-making, and institutions in quantitative system modelling". In: *Technological Forecasting and Social Change* 151, p. 119480.

Deffuant, Guillaume, Frédéric Amblard, Gérard Weisbuch, and Thierry Faure (2002). "How can extremism prevail? A study based on the relative agreement interaction model". In: *Journal of artificial societies and social simulation* 5.4.

Deffuant, Guillaume, David Neau, Frederic Amblard, and Gérard Weisbuch (2000). "Mixing beliefs among interacting agents". In: *Advances in Complex Systems* 3.01n04, pp. 87–98.

DeGroot, Morris H (1974). "Reaching a consensus". In: *Journal of the American Statistical Association* 69.345, pp. 118–121.

Drews, Stefan, Ivan Savin, and Jeroen van den Bergh (2022). "Biased perceptions of other people's attitudes to carbon taxation". In: *Energy Policy* 167, p. 113051.

Durham, William H (1991). *Coevolution: Genes, culture, and human diversity*. Stanford University Press.

Epstein, Joshua M and Robert Axtell (1996). *Growing artificial societies: social science from the bottom up*. Brookings Institution Press.

Fáth, Gábor and Miklos Sarvary (2005). “A renormalization group theory of cultural evolution”. In: *Physica A: Statistical Mechanics and its Applications* 348, pp. 611–629.

Fielding, Kelly S and Matthew J Hornsey (2016). “A social identity analysis of climate change and environmental attitudes and behaviors: Insights and opportunities”. In: *Frontiers in psychology* 7, p. 121.

Fisher, Ronald Aylmer and Frank Yates (1953). *Statistical tables for biological, agricultural and medical research*. Hafner Publishing Company.

Foramitti, Joël, Ivan Savin, and Jeroen van den Bergh (2021). “Regulation at the source? Comparing upstream and downstream climate policies”. In: *Technological Forecasting and Social Change* 172, p. 121060.

Geschke, Daniel, Jan Lorenz, and Peter Holtz (2019). “The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers”. In: *British Journal of Social Psychology* 58.1, pp. 129–149.

Grabisch, Michel and Agnieszka Rusinowska (2020). “A survey on nonstrategic models of opinion dynamics”. In: *Games* 11.4, p. 65.

Granovetter, Mark (1978). “Threshold models of collective behavior”. In: *American journal of sociology* 83.6, pp. 1420–1443.

Granovetter, Mark S (1973). “The strength of weak ties”. In: *American journal of sociology* 78.6, pp. 1360–1380.

Grimson, Alejandro (2010). “Culture and identity: two different notions”. In: *Social Identities* 16.1, pp. 61–77.

Gupta, Akhil and James Ferguson (1992). “Beyond” culture”: Space, identity, and the politics of difference”. In: *Cultural anthropology* 7.1, pp. 6–23.

Hamis, Sara, Stanislav Stratiev, and Gibin G Powathil (2021). “Uncertainty and sensitivity analyses methods for agent-based mathematical models: An introductory review”. In: *The Physics of Cancer: Research Advances*, pp. 1–37.

Hegselmann, Rainer, Ulrich Krause, et al. (2002). “Opinion dynamics and bounded confidence models, analysis, and simulation”. In: *Journal of artificial societies and social simulation* 5.3.

Helfmann, Luzie, Jobst Heitzig, Péter Koltai, Jürgen Kurths, and Christof Schütte (2021). “Statistical analysis of tipping pathways in agent-based models”. In: *The European Physical Journal Special Topics* 230.16, pp. 3249–3271.

Henrich, Joseph (2001). “Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change”. In: *American Anthropologist* 103.4, pp. 992–1013.

Henrich, Joseph and Richard McElreath (2003). “The evolution of cultural evolution”. In: *Evolutionary Anthropology: Issues, News, and Reviews: Issues, News, and Reviews* 12.3, pp. 123–135.

Herman, Jon and Will Usher (Jan. 2017). “SALib: An open-source Python library for Sensitivity Analysis”. In: *The Journal of Open Source Software* 2.9.

Kaaronen, Roope Oskari and Nikita Strelkovskii (2020). "Cultural evolution of sustainable behaviors: Pro-environmental tipping points in an agent-based model". In: *One Earth* 2.1, pp. 85–97.

Kapeller, Marie Lisa, Georg Jäger, and Manfred Füllsack (2019). "Homophily in networked agent-based models: a method to generate homophilic attribute distributions to improve upon random distribution approaches". In: *Computational Social Networks* 6.1, pp. 1–18.

Kashima, Yoshihisa (2008). "A social psychology of cultural dynamics: Examining how cultures are formed, maintained, and transformed". In: *Social and Personality Psychology Compass* 2.1, pp. 107–120.

Konc, Théo, Stefan Drews, Ivan Savin, and Jeroen van den Bergh (2022). "Co-dynamics of climate policy stringency and public support". In: *Global Environmental Change* 74, p. 102528.

Konc, Théo and Ivan Savin (2019). "Social reinforcement with weighted interactions". In: *Physical Review E* 100.2, p. 022305.

Kraan, Oscar, Steven Dalderop, Gert Jan Kramer, and Igor Nikolic (2019). "Jumping to a better world: An agent-based exploration of criticality in low-carbon energy transitions". In: *Energy Research & Social Science* 47, pp. 156–165.

Kuperman, Marcelo N (2006). "Cultural propagation on social networks". In: *Physical Review E* 73.4, p. 046139.

Lacasse, Katherine (2016). "Don't be satisfied, identify! Strengthening positive spillover by connecting pro-environmental behaviors to an "environmentalist" label". In: *Journal of Environmental Psychology* 48, pp. 149–158.

Laibson, David (1997). "Golden eggs and hyperbolic discounting". In: *The Quarterly Journal of Economics* 112.2, pp. 443–478.

Lamperti, Francesco, Giovanni Dosi, Mauro Napoletano, Andrea Roventini, and Alessandro Sapiò (2018). "Faraway, so close: Coupled climate and economic dynamics in an agent-based integrated assessment model". In: *Ecological Economics* 150, pp. 315–339.

Ligmann-Zielinska, Arika, Daniel B Kramer, Kendra Spence Cheruvilil, and Patricia A Soranno (2014). "Using uncertainty and sensitivity analyses in socioecological agent-based models to improve their analytical performance and policy relevance". In: *PLoS one* 9.10, e109779.

Loewenstein, George and Drazen Prelec (1992). "Anomalies in intertemporal choice: Evidence and an interpretation". In: *The Quarterly Journal of Economics* 107.2, pp. 573–597.

Lord, Charles G, Lee Ross, and Mark R Lepper (1979). "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence". In: *Journal of personality and social psychology* 37.11, p. 2098.

Mason, Winter A, Frederica R Conrey, and Eliot R Smith (2007). "Situating social influence processes: Dynamic, multidirectional flows of influence within social networks". In: *Personality and social psychology review* 11.3, pp. 279–300.

McQuarrie, Donald A (2000). *Statistical mechanics*. Sterling Publishing Company.

Mesoudi, Alex (2016). "Cultural evolution: a review of theory, findings and controversies". In: *Evolutionary biology* 43.4, pp. 481–497.

Niamir, Leila, Tatiana Filatova, Alexey Voinov, and Hans Bressers (2018). "Transition to low-carbon economy: Assessing cumulative impacts of individual behavioral changes". In: *Energy policy* 118, pp. 325–345.

Niamir, Leila, Gregor Kiesewetter, Fabian Wagner, Wolfgang Schöpp, Tatiana Filatova, Alexey Voinov, and Hans Bressers (2020). "Assessing the macroeconomic impacts of individual behavioral changes on carbon emissions". In: *Climatic change* 158.2, pp. 141–160.

Nickerson, Raymond S (1998). "Confirmation bias: A ubiquitous phenomenon in many guises". In: *Review of general psychology* 2.2, pp. 175–220.

Nigbur, Dennis, Evanthisia Lyons, and David Uzzell (2010). "Attitudes, norms, identity and environmental behaviour: Using an expanded theory of planned behaviour to predict participation in a kerbside recycling programme". In: *British journal of social psychology* 49.2, pp. 259–284.

Rai, Varun and Scott A Robinson (2015). "Agent-based modeling of energy technology adoption: Empirical integration of social, behavioral, economic, and environmental factors". In: *Environmental Modelling & Software* 70, pp. 163–177.

Railsback, Steven F and Volker Grimm (2019). *Agent-based and individual-based modeling: a practical introduction*. Princeton university press.

Safarzyńska, Karolina and Jeroen van den Bergh (2022). "ABM-IAM: optimal climate policy under bounded rationality and multiple inequalities". In: *Environmental Research Letters* 17.9, p. 094022.

Saltelli, Andrea (2002). "Making best use of model evaluations to compute sensitivity indices". In: *Computer physics communications* 145.2, pp. 280–297.

Savin, Ivan, Felix Creutzig, Tatiana Filatova, Joël Foramitti, Théo Konc, Leila Niamir, Karolina Safarzynska, and Jeroen van den Bergh (2022). "Agent-based modeling to integrate elements from different disciplines for ambitious climate policy". In: *Wiley Interdisciplinary Reviews: Climate Change*, e811.

Schaller, Mark and Michael Muthukrishna (2021). "Modeling cultural change: Computational models of interpersonal influence dynamics can yield new insights about how cultures change, which cultures change more rapidly than others, and why." In: *American Psychologist* 76.6, p. 1027.

Schelling, Thomas C (1969). "Models of segregation". In: *The American economic review* 59.2, pp. 488–493.

Schultz, P Wesley, Jessica M Nolan, Robert B Cialdini, Noah J Goldstein, and Vladas Griskevicius (2007). "The constructive, destructive, and reconstructive power of social norms". In: *Psychological science* 18.5, pp. 429–434.

Smaldino, Paul E (2019). "Social identity and cooperation in cultural evolution". In: *Behavioural processes* 161, pp. 108–116.

Sobol, Ilya M (2001). "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates". In: *Mathematics and computers in simulation* 55.1-3, pp. 271–280.

Van der Werff, Ellen, Linda Steg, and Kees Keizer (2013). "It is a moral issue: The relationship between environmental self-identity, obligation-based intrinsic motivation and pro-environmental behaviour". In: *Global environmental change* 23.5, pp. 1258–1265.

Van der Werff, Ellen, Linda Steg, and Kees Keizer (2014). "I am what I am, by looking past the present: the influence of biospheric values and past behavior on environmental self-identity". In: *Environment and behavior* 46.5, pp. 626–657.

Waring, Timothy M, Michelle Kline Ann, Jeremy S Brooks, Sandra H Goff, John Gowdy, Marco A Janssen, Paul E Smaldino, and Jennifer Jacquet (2015). "A multilevel evolutionary framework for sustainability analysis". In: *Ecology and Society* 20.2.

Watts, Duncan J and Steven H Strogatz (1998). “Collective dynamics of ‘small-world’ networks”. In: *Nature* 393.6684, pp. 440–442.

Wynes, Seth and Kimberly A Nicholas (2017). “The climate mitigation gap: education and government recommendations miss the most effective individual actions”. In: *Environmental Research Letters* 12.7, p. 074024.

Yi, Richard, Kirstin M Gatchalian, and Warren K Bickel (2006). “Discounting of past outcomes.” In: *Experimental and clinical psychopharmacology* 14.3, p. 311.

Appendix

2.A Additional figures




Figure 2.A.1: *The beta distribution used to generate initial values for individuals' attitudes and thresholds for a variety of input a and b values. Only if a and b are equal in value will the distribution be symmetric. The ratio of a and b dictates the distribution mean $E(X) = a/(a + b)$, whilst smaller values of a and b lead to greater initial identity polarisation.*




Figure 2.A.2: *The initial distribution of individuals' attitudes correlates strongly with the total societal emissions per agent and behaviour. The contour borders are approximately linear representing lines of a constant expectation value of the Beta distribution $a_A/(a_A + b_A)$. This contour plot is produced from 20480 experiments including 4096 combinations of parameters a_A and b_A .*




Figure 2.A.3: *Within "identity bubbles" there is uniform weighting amongst agents. The left figure shows an identity time series and the formation of "identity bubbles", corresponding to the different colours. The right-hand sub-figure shows the mean social network weighting $\alpha_{n,k}$ between members within an "identity bubble" as a solid line. The dashed lines are $1/N_C$, where N_C is the number of members within the cluster. The mean cluster weighting in each group tends to $1/N_C$ when clusters are well separated, as can be seen in the middle of the simulation run. This run corresponds to the dynamic weighting in Figure 2.5.*




Figure 2.A.4: *Sobol sensitivity analysis of model parameters with the inclusion of green influencers where we test a range of [1,100] agents for a total of 1920 experiments.*

Chapter 3

The cultural multiplier of climate policy*

3.1 Introduction

To achieve the deep decarbonisation required to meet emissions targets, consumption changes are needed across the board, i.e. applying to all goods, services and hence production sectors. This may involve a shift in lifestyles (Girod et al., 2014). As part of this, there can be positive and negative spillovers (Truelove et al., 2014; Lanzini and Thøgersen, 2014) between distinct consumption categories, such as travel by plane and following a vegan diet (Andersson and Nässén, 2022). To better understand the mechanism and implications, we develop a model that connects a multiple consumption category module to a social network module and assess how a carbon tax induces decarbonisation. To study large-scale shifts in consumption behaviours due to repeated and cumulative social interactions, we conduct our analysis through the lens of cultural change (Davis et al., 2018; Kaaronen and Strelkovskii, 2020; Sovacool and Griffiths, 2020). Building on computational models of cultural change (Epstein and Axtell, 1996; Axelrod, 1997; Kuiperman, 2006; Torren-Peraire et al., 2024), we study how long-term preference dynamics may affect decarbonisation.

This paper contributes to the limited literature on carbon taxation with endogenous preference changes. The model proposed in Mattauch et al. (2022) studies the impact of carbon pricing which directly affects preferences for clean goods. This is described as a crowding-in (greater consumption of clean good) or -out (greater consumption of dirty good) effect of carbon pricing. Konc, Savin, and Bergh (2021) study the role of direct and indirect effects of carbon taxation, through price and social influence mechanisms. Their focus centres on the impact of the social interactions on climate policy effectiveness, defined as a “social multiplier” (Glaeser et al., 2003; Konc, Savin, and Bergh, 2021). In this article, we extend this concept to the case of social learning across multiple consumption categories (e.g. transport, food, tourism), labelling this a “cultural multiplier”, to assess how social influence affects lifestyle changes. In contrast to previous economics literature on the interaction of culture and the environment (Schumacher, 2015; Bezin, 2019), we choose not to take a generation-based modelling approach due to the limited timeframe over which decarbonisation must occur, e.g. EU net-zero by 2050 target. However, in a similar fashion to Schumacher (2015) we model how the emergence of pro-environmental

*This chapter has been submitted and is under revision.

culture can stimulate greater decarbonisation in a positive feedback loop.

Low-carbon consumption choices are influenced by a multitude of factors, not just pricing (Wang et al., 2021). Social influences such as injunctive and descriptive norms (Davis et al., 2018), the framing of a carbon tax (Hartmann et al., 2023) or similarities in low- and high-carbon alternatives, can play significant roles. For example, individuals may perceive low-carbon goods such as electric vehicles (EV) to be less substitutable because of different technological characteristics, seen in phenomena such as range anxiety (Pevec et al., 2020). As a consequence, urban communities can have higher rates of EV adoption than rural ones (Westin et al., 2018).

A key component in determining the dynamics of socially informed preferences is the structure of the network within which social interactions occur. Individuals with high socioeconomic status (Nielsen et al., 2021), who occupy central positions in social networks, can generate social tipping points for decarbonisation. However, if their preferences are not consistently aligned with decarbonisation goals they may inhibit a rapid transition to low-carbon alternatives through their sustained high-carbon consumption (Mattioli et al., 2023).

Our central research question is: How does cultural change moderates the effectiveness of carbon taxes. In pursuit of this goal, we explore the following sub-questions: Is the cultural multiplier similar in size to the social multiplier? And what are the socioeconomic characteristics affecting the magnitude of the cultural multiplier? Regarding the latter, we will consider substitutability between low- and high-carbon goods, social network structure, proximity of like-minded individuals in a social network, and diversity of lifestyles.

The remainder of the paper is organised as follows. In Section 3.2 we formulate the model of market processes, social network interactions, cultural change, and climate policy. In Section 3.3 we analyse the strength of the cultural multiplier, compare it to the social multiplier, and explore which socioeconomic factors shape its size. Section 3.4 concludes.

3.2 The model

3.2.1 Conceptual approach

We construct a model of individuals' consumption behaviour subject to a carbon tax. To capture the interaction of endogenous preference driven by social interactions and climate policy we use an Agent-Based Model (ABM). Individuals act as utility-maximising agents with heterogeneous preferences for low-carbon goods that evolve through social imitation.

Figure 1 provides a schematic overview of the model. The yellow boxes represent the consumption choice of individuals and are constructed from the utility function subject to a budget constraint. The stacked yellow boxes indicate that consumption occurs in multiple sectors, such as energy, transport or food. The blue boxes represent the social imitation module of the model which produces dynamic preferences for low-carbon consumption. This imitation occurs via weighted interaction in a social network (Konc and Savin, 2019). These interactions are mediated in strength by the similarity of individuals in their environmental identity. The grey stacked box indicates the multiple individuals that compose the social network. Individuals' consumption choices produce emissions each time step which contribute to network-wide cumulative emissions, shown in the red boxes. Lastly, the green boxes are the climate policy module, capturing the role of carbon

taxation in guiding the social network towards low-carbon consumption. Specifically, a tax applied to consumption across all the different categories influences high-carbon good prices and produces revenues which are redistributed as a lump sum.




Figure 3.2.1: *Model structure*. Blue boxes capture the social interactions; yellow the preference dynamics and consumption choices; green the climate policy; and red the carbon emissions. Arrows indicate the direction of influence. Grey stacked boxes indicate multiple individuals whilst coloured stacked boxes show multiple consumption categories.

3.2.2 Market module

The model employs a constant elasticity of substitution (CES) utility function where individuals have preferences for low-carbon and high-carbon alternatives. Highly substitutable goods such as energy, where the source has little impact on utility, have strongly non-linear responses to price changes or a given preference between alternatives. However, to answer the main research question of how carbon taxation is affected by cultural change requires representing consumption habits across multiple human needs, i.e. lifestyle (Foramitti et al., 2024). To capture this, we extend the CES function to multiple consumption categories which results in a nested CES utility function (NCES). Individuals' make choices of consumption quantities between low- L and high-carbon H goods across multiple consumption categories m . In the model at each discrete time step t , individuals, denoted by index i , maximize their utility based on their preference for low-carbon goods, $A_{t,i,m}$, subject to a budget constraint imposed by individuals' expenditure B ,

$$\max_{L_1, \dots, L_M, H_1, \dots, H_M} U(L_1, \dots, L_M, H_1, \dots, H_M, a_1, \dots, a_M, A_1, \dots, A_M, \sigma_1, \dots, \sigma, \nu) \quad (3.1)$$

The utility is given by a NCES function with two levels (Sato, 1967)

$$U_{t,i} = \left(\sum_{m=1}^M a_m \left(A_{t,i,m} L_{t,i,m}^{\frac{\sigma-1}{\sigma}} + (1 - A_{t,i,m}) H_{t,i,m}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma(\nu-1)}{(\sigma-1)\nu}} \right)^{\frac{\nu}{\nu-1}}, \quad (3.2)$$

At the top level, we represent different consumption categories, while on the bottom level (within each category) there are low- and high-carbon goods alternatives (Konec,

Savin, and Bergh, 2021; Mattauch et al., 2022), $L_{t,i,m}$ and $H_{t,i,m}$ and substitutability between each of the two alternatives σ . Between categories, there is a further preference parameter for consumption a_m where $\sum_{m=1}^M a_m = 1$ and substitutability across categories ν . To reduce model complexity, for each individual i we assume the preference parameter between categories a_m , e.g. the preference for transport over energy consumption, is fixed and there is equal weighting of consumption categories $\sum_{m=1}^M a_m = 1$. Utility maximisation is subject to a budget constraint

$$B = \sum_{m=1}^M L_{t,i,m} P_{L,t,m} + H_{t,i,m} P_{H,m}. \quad (3.3)$$

From the Lagrangian first-order conditions for the system we derive demand relationships for the low- and high-carbon alternatives for all consumption categories as functions of preferences, prices and degrees of substitution (see 3.A). The demands for high- and low-carbon goods are given by

$$H_{t,i,m} = \frac{B\chi_{t,i,m}}{Z_{t,i}} \quad (3.4)$$

$$L_{t,i,m} = \frac{B\Omega_{t,i,m}\chi_{t,i,m}}{Z_{t,i}} \quad (3.5)$$

Here $\Omega_{t,i,m}$ is the ratio between low- and high-carbon good quantities, $\chi_{t,i,m}$ captures the relative weighting of consumption between different categories, and $Z_{t,i}$ is a normalisation term. They are defined as:

$$\Omega_{t,i,m} = \frac{L_{t,i,m}}{H_{t,i,m}} = \left(\frac{P_{H,m} A_{t,i,m}}{P_{L,m} (1 - A_{t,i,m})} \right)^\sigma \quad (3.6)$$

$$\chi_{t,i,m} = \left(\frac{a_m A_{t,i,m}}{P_{L,m} \Omega_{t,i,m}^{\frac{1}{\sigma}}} \right)^\nu \left(A_{t,i,m} \Omega_{t,i,m}^{\frac{\sigma-1}{\sigma}} + 1 - A_{t,i,m} \right)^{\frac{\nu-\sigma}{(\sigma-1)}} \quad (3.7)$$

$$Z_{t,i} = \sum_{p=1}^M \chi_{t,i,p} (\Omega_{t,i,p} P_{L,t,p} + P_{H,t,p}) \quad (3.8)$$

3.2.3 Social imitation module

This module is closely adapted from the one proposed in Torren-Peraire et al. (2024), which in turn is derived from DeGroot (1974). Individual preferences for low-carbon consumption evolve due to social imitation of neighbours' consumption behaviour. A social network is introduced to represent the context within which this process occurs. The network is composed of N individuals i , each with ego networks N_i , which interact with each other each time-step t . Future preferences are a weighted average of current preferences $A_{t,i,m}$ and an external social imitation influence of others' low-carbon consumption behaviour $C_{t,i,m}$:

$$A_{t+1,i,m} = (1 - \phi_m) A_{t,i,m} + \phi_m \sum_{j=1}^{N_i} \alpha_{t,i,j}^{CM} C_{t,j,m}, \quad (3.9)$$

How sensitive an individual's preferences are to social influence is mediated by the social susceptibility parameter $\phi \in [0, 1]$. The parameter $\alpha_{t,i,j}^{CM}$ captures how much individual i values the opinion of neighbouring individual j . The initial preferences for

low-carbon consumption across multiple consumption categories $A_{0,i,m}$ are generated using a Beta distribution. We set the similarity between an individual's initial preferences for low-carbon consumption across different categories using an initial coherence parameter, $c \in [0, 1]$, where the case of $c = 0$ represents an individual whose preferences do not align with each other. Following Konc, Savin, and Bergh (2021), individuals copy the proportion of neighbours consumption that is low-carbon:

$$C_{t,j,m} = \frac{L_{t,j,m}}{L_{t,j,m} + H_{t,j,m}}. \quad (3.10)$$

For our model of the cultural multiplier, we assume that the preferences of individuals for low-carbon goods are not observable, instead, these must be inferred through observing actual consumption. The social multiplier of climate policy may be characterised as a misinterpretation of consumption change (from increased price) as a preference change. The result of this misinterpretation is that for a given carbon emission target a lower carbon tax is required. Alternatively, the actual decarbonisation effect exceeds the one predicted by fixed preference models of Pigouvian taxes, as the policy induces preference change (Koessler, Engel, et al., 2021). Given the importance of imitation, in 3.B we study how greater substitutability between goods leads to a more non-linear relationship between preferences and consumption shares.

In the model, the consumption behaviours of neighbours in the social network are not taken into consideration equally leading to a lack of global preference convergence (Dandekar et al., 2013). Instead, individuals strive for greater homophily through weighted social imitation. To model this we follow Axelrod (1997), where past interactions between pairs of individuals leads to stronger future interactions. The strength of these interactions, $\alpha_{t,i,j}^{CM}$, are dictated by the similarity in environmental identity $I_{t,i}$ (the mean of low-carbon preferences).

The social network weighting matrix $\alpha_{t,i,j}^{CM}$ introduced in Equation 3.9 is given by the softmax function (Konc and Savin, 2019)

$$\alpha_{t,i,j}^{CM} = \frac{e^{-\theta|I_{t,i} - I_{t,j}|}}{\sum_{j \neq i}^{N_i} e^{-\theta|I_{t,i} - I_{t,j}|}}, \quad (3.11)$$

Based on the Euclidean identity distance of N_i neighbours, with θ a measure of confirmation bias. The greater this bias, the less open individuals are to imitating the behaviour of neighbours with a different environmental identity. A simplified version of the identity model developed in Torren-Peraire et al. (2024) is used here. It defines the environmental identity of an individual as the average of their preferences for different categories of low-carbon goods. This model provides an indirect mechanism for spillovers, through which a greater pro-environmental identity can make green behaviours more likely (Van der Werff et al., 2014).

$$I_{t,i} = \frac{1}{M} \sum_{m=1}^M A_{t,i,m}. \quad (3.12)$$

A shift towards pro-environmental identities not only requires a change in one category of consumption, such as the growing popularity of a vegan diet, but coordination across multiple consumption categories in a low-carbon direction. This results in preference change producing a slower, longer-term cultural change. In order to compare the strength

of the social and cultural multipliers, a second weighting matrix $\alpha_{t,i,j,m}^{SM}$, is defined where individuals consider the Euclidean distance between preferences instead of identities,

$$\alpha_{t,i,j,m}^{SM} = \frac{e^{-\theta|A_{t,i,m} - A_{t,j,m}|}}{\sum_{\substack{j \neq i \\ N_i}} e^{-\theta|A_{t,i,m} - A_{t,j,m}|}}. \quad (3.13)$$

This characterises the social multiplier case, where individuals evaluate their similarities with neighbours separately for each of their m consumption categories. In summary, Equation 3.13 replaces Equation 3.11 in defining the social weighting term used in Equation 3.9. This can result in individuals paying varying amounts of attention to a neighbour's consumption behaviour depending on the consumption category.

To highlight the role of environmental identity as a mechanism for behavioural spillovers in consumption, consider two individuals: Alice and Bob. Each evaluates how much attention is paid to the others' consumption behaviour across three consumption categories: food, transport, and energy. Alice is a strict vegan and thus has a strong low-carbon preference for the consumption category food ($A_1 = 1$) but is relatively indifferent towards low-carbon consumption of other categories (e.g., $A_2 = 0.5, A_3 = 0.5$). Bob, on the other hand, is slightly indifferent towards the carbon intensity of all consumption categories (e.g., $A_1, A_2, A_3 = 0.45$). Let us first consider the case of the social multiplier, where individuals consider the weighting of social influence separately for each consumption category based on similarity in preferences. Bob may pay attention to Alice's behaviour in transport and energy consumption. However, in the case of food consumption, Bob sees Alice as an outlier or radical, choosing to ignore her behaviour. In the case of the cultural multiplier, things work differently. Here Bob considers the similarity in identity between themselves and Alice, considering their proximity in preferences across multiple consumption categories (see Equation 3.12). Relative to the social multiplier case in the cultural multiplier, Bob pays less attention to Alice in transport and energy categories, but crucially does not ignore their behaviour in the food category. Over repeated social interactions, Alice influences Bob towards a more low-carbon preference in food which would have been ignored under the social multiplier scenario. In this fashion, selective imitation based on environmental identity at the individual level acts as a mechanism for generating cohesion in consumption across the entire social network.

Social influence between individuals is facilitated in the model via a social network. The specification of the network can be adapted to capture the relevant context in terms of consumption category (e.g., more or less conspicuous) and medium of social interaction (e.g., face-to-face, word of mouth, online, or geographical). We consider three different network structures, each with heterogeneous degree distribution: Small-world (SW (Watts and Strogatz, 1998)), Stochastic Block Model (SBM (Holland et al., 1983)) and scale-free (SF (Albert and Barabási, 2002)) - see Figure 3.2.2.

The SW model represents physical (offline) social networks. The small-world property is generated through high clustering of nodes with short path lengths due to a few long-range weak ties across the network. The SBM allows for the representation of clustered groups of nodes that have higher connection density within blocks than between them. In the model, we consider how the dichotomous relationship between two blocks can affect decarbonisation across the entire network. This network structure facilitates the study of consumption decarbonisation in loosely linked communities, such as rural versus urban settings. For the SF network, the use of a growing network with preferential attachment generates a degree distribution that follows a power law. This results in a handful of

nodes having a high number of connections whilst most have few, such as in online social networks. This network structure may be used to study the role of individuals with high socioeconomic status (Nielsen et al., 2021). Specifically of interest is how a central hegemony of low- or high-carbon consumption in the highest degree nodes may tip the rest of the system due to their far-reaching social influence. Lastly, we introduce a measure of homophily $h \in [0, 1]$, which indicates the distribution of initial environmental identities among neighbours. For $h = 0$, individuals are randomly positioned within the network, while for $h = 1$ individuals connected have the smallest possible distance in environmental identities.




Figure 3.2.2: *Illustrative example of different social network structures tested.*

3.2.4 Climate policy module

Without loss of generality, we assume that high-carbon goods and services have emissions of one and the low-carbon good have zero emissions. The consumption of high-carbon goods by individuals produces emissions which contribute to a global cumulative quantity E , according to

$$E = \sum_{t=0}^{t_{max}} \sum_{i=1}^N \sum_{m=1}^M H_{t,i,m}. \quad (3.14)$$

where there are t_{max} time steps in each experiment. A carbon tax, τ , is introduced resulting in a higher price for high-carbon goods

$$P_{H,m} = P_{B,H,m} + \tau \quad (3.15)$$

The price of low-carbon goods and base price of high-carbon goods are $P_{L,m}$, $P_{B,H,m} = 1$. The revenues of the carbon tax are recycled to consumers using a lump-sum carbon dividend. The analysis for a carbon tax can also be applied to a more general carbon price. However, to better integrate our model with current literature on endogenous preferences we choose to label the climate policy intervention as a carbon tax (Konc, Savin, and Bergh, 2021; Mattauch et al., 2022).

3.3 Results

3.3.1 Overview of experimental runs

Each experimental run consists of 3000 individuals interacting over 360 time steps. This can be considered to represent 30 years with each time step being a month. Unless stated

otherwise in the following results each sub-figure is composed of 30,000 experimental runs. Individual initial preferences are drawn from a Beta distribution symmetrical about indifference towards the carbon content of goods ($\bar{A}_m = 0.5$). Appendix Figure 3.C.1 gives an illustrative example of a typical model run showing the environmental identity trajectories for zero and low-carbon tax of $\tau = 0.15$. The parameters used in models are shown in Table 3.3.1 and a summary of variables in Table 3.3.2. Note that Network parameters such as number of edges to attach from new nodes in the SF and SW, or the intra- and inter-block densities are configured to ensure that across the three networks tested the density of connection is 0.1.

Table 3.3.1: Model Parameters

Parameter Name	Symbol	Description	Default Value	Range Tested
Initial preference Beta	a, b	Beta distribution parameters used to generate initial preferences	2	[0.1, 8]
Low- and high-carbon substitutability	σ	Elasticity of substitution between low-carbon and high-carbon goods	4	[1.1, 8]
Between category substitutability	ν	Elasticity of substitution across consumption categories	2	[1.1, 8]
Social susceptibility	ϕ	Influence of social imitation on preference for low-carbon consumption	0.02	[0,1]
Confirmation bias	θ	Confirmation bias towards individuals with similar environmental identities	5	[0, 50]
Carbon tax	τ	Tax imposed on high-carbon goods	[0,1]	[0, 5]
Consumption categories	M	Number of categories of goods	2	[1, 50]
Total individuals	N	Total number of individuals in the model	3000	[500,3000]
Homophily state	h	Degree of initial similarity between neighbours in terms of environmental identity	0	[0, 1]
Coherence state	c	Similarity individual's low-carbon preferences across consumption categories.	0.9	[0, 1]
Maximum time steps	t_{max}	Total time-steps in experimental runs	360	
Price of low-carbon goods	$P_{L,m}$	Price of low-carbon goods in category m	1	
Base price of high-carbon good	$P_{B,H,m}$	Base price of high-carbon goods in category m	1	
SF density		Density of connections between individuals	0.1	
SBM block number		Number of blocks in stochastic block model	2	
SBM intra-block density		Density of connection between individuals within block	0.02	
SBM inter-block density		Density of connection between individuals between blocks	0.005	
SW density		Density of connections between individuals	0.1	
SW probability rewire		Probability of rewiring to produce long distance ties	0.1	
Stochastic Seed Repetitions		Variations of initials seed for preferences, network structure, homophily and coherence	100	

Table 3.3.2: Model Variables

Variable Name	Symbol	Description
Low-carbon consumption	$L_{t,i,m}$	Quantity of low-carbon goods consumed by individual i in category m at time t
High-carbon consumption	$H_{t,i,m}$	Quantity of high-carbon goods consumed by individual i in category m at time t
Price of high-carbon goods	$P_{H,m}$	Price of high-carbon goods in category m , sum of base price and carbon tax
Preference for low-carbon goods	$A_{t,i,m}$	Preference for low-carbon goods by individual i in category m at time t
Environmental identity	$I_{t,i}$	Average low-carbon preference of individual i across all categories at time t
Social influence weighting	$\alpha_{t,i,j}$	Weight assigned by individual i to the consumption behaviour of neighbour j
Ratio of low to high carbon goods	$\Omega_{t,i,m}$	Ratio between low-carbon and high-carbon consumption by individual
Consumption weighting factor	$\chi_{t,i,m}$	Weighting of consumption across categories
Normalisation term	$Z_{t,i}$	Normalisation term used in demand functions
Expenditure	B	Budget available for consumption, $1/N$
Preference for categories	a_m	Preference for consumption categories, $1/M$

3.3.2 Evaluating the cultural multiplier

To assess the impact of cultural change on the effectiveness of climate policy we measure cumulative carbon emissions E under three conditions:

1. Fixed preferences – no social influence of preferences and consumption decisions.
2. Social multiplier - dynamic preferences due to social imitation through preference similarity for each consumption category separately, as captured by Equation 3.13
3. Cultural multiplier - dynamic preferences due to social imitation through identity similarity, as captured by Equation 3.11

In the model, social influence occurs through the imitation of consumption behaviours which depend on both preferences and the carbon tax level. To test whether the cultural multiplier is a function of the strength of the carbon tax we consider a range of values $\tau = [0, 1]$. Figure 3.3.1 shows the cumulative emissions for the three different cases in a small-world network. For each carbon tax value and network structure, we ran 100 experiments with different stochastic seeds for the initial preference distribution, network structure, distribution of individuals in the network and coherence in preferences. In the figure the solid line shows the mean and shaded region showing the 95% confidence interval over stochastic seed runs.

In the case of the cultural multiplier (green) we see a large and instant decrease in cumulative emissions with the introduction of a carbon tax, relative to the fixed preferences scenario. The imitation of consumption choices results in additional decarbonisation across multiple consumption categories. When comparing the social (orange) and cultural multiplier cases (green) we find that the latter has much lower emissions for small carbon taxes (approximately $\tau < 0.5$). This greater strength of the cultural multiplier

over the social multiplier can be explained by the consensus-forming effect of environmental identity. When evaluating individuals across multiple preferences collectively it becomes harder for outliers in the preference space to isolate themselves into communities with a high degree of preference homophily. This results in individuals imitating a wider range of consumption behaviours, and reach faster consensus formation, translating into greater emissions reduction of carbon taxation.




Figure 3.3.1: *Cumulative emissions for the case of fixed preferences (blue), cultural multiplier (green) and social multiplier (orange) using a small-world network. The shaded region indicated 95% confidence interval for 100 stochastic runs. The Stochastic Block Model and scale-free graphs are shown in Appendix Figure 3.C.2.*

To compare the findings with previous work by Konc, Savin, and Bergh (2021) we study the tax reduction M_{tax} , induced by the social and carbon tax multipliers, defined as

$$M_{tax} = 1 - \frac{\tau_s}{\tau_f} \quad (3.16)$$

where τ_f is the carbon tax required in the fixed preferences case to match the emissions reduction caused by a carbon tax τ_s in the social and cultural multiplier cases. Using data from Figure 3.3.2 on cumulative emissions we can map the required τ_f value onto a τ_s value i.e. for a given emissions target what is the carbon tax required in the fixed preferences, social and cultural multiplier cases. However, given that cumulative emissions are much lower with the social and cultural multipliers than in the fixed preferences case, additional simulations are needed to determine the carbon tax required for the fixed preferences case to match the emissions of the other two across all carbon tax levels. With this aim, we first calculate what the maximum and minimum cumulative emission produced by the social and cultural multiplier cases are for carbon tax values $\tau_s = [0, 1]$. Secondly, we use these extreme emission values as targets and calculate what the required carbon tax τ_f would be to achieve this in the fixed preferences case.

Figure 3.3.2 shows that the cultural multiplier has a greater tax reduction effect than the social multiplier. Additionally, due to the repeated nature of the interactions in the model, the mean magnitude of the emissions reduction ($M_{tax} > 0.5$) of both multipliers

is much stronger than that identified in Konc, Savin, and Bergh (2021) ($M_{tax} = 0.38$). With each time step, a growing share of an individual's current preferences are shaped cumulatively by the behaviour of their neighbours, potentially away from their initial preferences. For high-carbon taxes, $\tau > 0.5$ the difference between the social and cultural multiplier cases almost vanishes. At these values, the strength of the carbon tax signal overwhelms any nuances in how the preference change due to social influence occurs. Even the most stubborn of individuals with strong high-carbon preferences choose to pursue low-carbon consumption in all categories, with social imitation accelerating this change in lifestyles. In other words, if high-carbon goods are sufficiently expensive relative to the low-carbon alternative, then the specifics of who individuals choose to imitate no longer matter; the network tips collectively towards low-carbon consumption. In the case of very small carbon taxes $\tau < 0.05$, there is a much greater variance in the tax reduction M_{tax} due to the strong path dependency of the model. Small differences in social interactions can lead to radically different preference outcomes and, consequently, emissions due to the polarizing effect of consumption imitation. As opposed, the presence of a greater carbon tax steers the system towards a narrower set of equilibrium states.




Figure 3.3.2: *Carbon tax reduction for social multiplier (orange) and cultural multiplier (green). A small-world network is used; results for alternative networks are shown in Appendix Figure 3.C.3. The shaded region indicated 95% confidence interval for 100 stochastic seed runs.*

3.3.3 Impact of key parameters on the cultural multiplier

Impact of substitutability between low- and high-carbon goods on the cultural multiplier

To better understand how sensitive the cultural multiplier is to different socio-economic conditions, we vary the latter and assess how this influences the carbon tax's effectiveness. Improvements in low-carbon technologies such as plant-based alternatives, better EV charging facilities or investment in public transport may increase the substitutability between low- and high-carbon goods. Therefore, different levels of low- and high-carbon

good substitutability σ can represent degrees of low-carbon technological progress or infrastructure availability. We now consider the strength of the cultural multiplier under these varying scenarios. In Figure 3.3.3 we plot cumulative emissions for the cultural multiplier case (preference spillovers) for different substitutability σ in the SW network (see Appendix Figure 3.C.4 for SBM and SF networks). For high-carbon tax rates, the greater the substitutability, the greater the decarbonisation. This occurs because individuals receive a much lower penalty in the utility function for concentrating their consumption in one good, allowing them to better exploit the price asymmetry between low- and high-carbon alternatives. However, this same concentration of consumption results in polarised consumption proportions $C_{t,i,m}$. Individuals then imitate these consumption proportions, gradually leading to polarisation in low-carbon preferences. This can inhibit the spread of low-carbon consumption as individuals who have high-carbon preferences are able to express this preference in their consumption. Consequently, high-carbon groups of individuals isolate themselves by avoiding interactions with “greener” neighbours. This effect hinders decarbonisation at low-carbon tax levels, reversing the emissions pattern across substitutability scenarios. In contrast to these dynamics in the extreme case of $\sigma = 1.01$, the emissions curves detach from that of larger substitutability values. Here, individuals are unable to fully express their preferences in their consumption, and thus, social imitation becomes less representative of true beliefs, resulting in a lower carbon tax effectiveness.




Figure 3.3.3: *Impact of substitutability between low- and high-carbon alternatives on cumulative emissions*

To further study the relationship between the social and cultural multiplier we compare the two for different levels of substitutability σ , see Figure Figure 3.3.4. Note that in this figure we only consider the small-world network case. For very low substitutability $\sigma = 1.01$, the social and cultural multiplier converge. As previously identified in Figure 3.3.3, this low substitutability causes a breakdown in effective social learning through imitation. This results in less opportunity for behavioural spillovers to be leveraged, hence the small differences in emissions between the social and cultural multiplier. On the other hand, increasing substitutability enhances the cultural multiplier, whereas the effect of the social multiplier shows a comparatively smaller gain. This effect is especially pertinent at low tax values where the gap between the two multipliers is largest. The

emissions curve for the social multiplier (orange) shifts towards the fixed preferences case (blue) for increasing substitutability (across sub-figures). As highlighted in Figure 3.3.3, a greater σ parameter facilitates the concentration of consumption into either a low or high-carbon good. This polarisation enables the formation of social bubbles in which individuals who prefer high-carbon goods only communicate with like-minded individuals. These effects make preference change more challenging in the social multiplier case, as individuals become less responsive to the consumption habits of others outside their social bubble. Therefore, the absence of a behavioural spillover mechanism in the social multiplier case results in a lock-in of high-carbon consumption for individuals with a preference for high-carbon goods ($A < 0.5$). Consequently, in the social multiplier case at large substitutabilities, a larger part of consumption changes occur primarily due to carbon tax increases across experiments. However, a much higher substitutability does not significantly increase the gap between the social and cultural multiplier. This can be seen in comparing Figure 3.3.1, $\sigma = 4$, with the $\sigma = 10$ experiments in Figure 3.3.3.




Figure 3.3.4: *Differences between social and cultural multiplier at different substitutabilities between low- and high-carbon goods.*

Global sensitivity analysis

To study how the relative importance of model parameters and bolster the robustness of our previous results we conduct Sobol sensitivity analysis (Sobol, 2001). This is implemented with the SALib python library (Herman and Usher, 2017) for the cases of a SW, SBM and SF networks. Parameter ranges tested and fixed parameters are shown Table 3.3.1. For each of the 11 variables, we take 128 values with a mean of 20 stochastic variations for a total of 184,320 experiments. The total order sensitivity of final cumulative emissions is shown in Figure 3.3.5, with the first order index depicted in Appendix Figure 3.C.5. Cumulative emissions are primarily determined by the initial preferences Beta a parameter. A greater value relative to the Beta b parameter results in an initial preference distribution which results in more low-carbon consumption. The carbon tax also strongly contributes to the cumulative emissions variance as these induce low-carbon consumption through price inequalities. Additionally, in the Total Sobol index, we see that the low-carbon substitutability is significant. This is due to greater non-linearity in consumption choices with greater substitutability as described in 3.B. Lastly, in Figure 3.3.5, we do not see variation across the different network structures in the importance

of parameters.




Figure 3.3.5: *Total Sobol sensitivity analysis for 128 parameter values for a total of 184,320 experiments*

Homophily, network structure and lifestyle diversity

When comparing the social and cultural multiplier cumulative emissions curves over a range of carbon taxes for the SW, SBM and SF networks, we see little variation in results (shown in Figure 3.3.1 and Appendix Figure 3.C.2). This can be explained by the fact that, in these experimental runs, we assume no homophily in the initial environmental identity of individuals within the network. However, ideological polarisation surrounding climate change is growing (Falkenberg et al., 2022). Therefore, a key avenue for study is how this polarisation of preferences within social networks, and the structure of networks themselves, can inhibit decarbonisation of consumption (Flüe and Vogt, 2024). In our model we represent this polarisation by considering initial homophily in environmental identity: The extent to which individuals are surrounded by like-minded neighbours. Across the SW, SBM and SF networks we find that greater homophily in initial preferences between neighbours leads to a higher effective carbon tax rate to meet the same emissions reduction - see 3.C.1 for an in-depth study.

Given the importance of behavioural spillover shown in the cultural versus social multiplier, we now consider how a richer representation of lifestyles can affect decarbonisation by increasing the complexity of consumption decisions. In Appendix Figure 3.C.6, we, therefore, vary the number of consumption categories modelled, M , for three different carbon tax values. For the cultural multiplier, at low carbon tax values, the addition of more consumption categories leads to lower emissions, relative to the social multiplier case. The greater M , the lower the impact of extreme preferences on the formation of environmental identity I . This results in greater consensus formation and thus a faster collective shift to low-carbon lifestyles. However, at high carbon tax values, as in Figure 3.3.1, there is no distinction between the cultural and social multiplier due to the constraining force of the price signal on consumption choices.

3.4 Conclusions

This study highlights the joint effect of carbon taxes and cultural change in fostering low-carbon consumption. Our model builds on previous literature on carbon taxes with endogenous preference change but focuses on a longer-term cultural change. We extend the concept of a social multiplier of environmental policy (increased emissions reduction from a carbon tax due to social imitation (Konc, Savin, and Bergh, 2021) to the case of repeated social imitation in multiple consumption categories and define this as the cultural multiplier. In our novel agent-based model, individuals make consumption choices between low- and high-carbon goods across multiple categories. These individuals have heterogeneous and dynamic preferences for low-carbon goods, which evolve through repeated and weighted social interactions. The model assesses how cultural change may enhance or hinder the impact of carbon taxes. Additionally, we identify which socio-economic characteristics influence the magnitude of the cultural multiplier.

Our results show that incorporating change in endogenous preferences through the cultural multiplier significantly enhances the effectiveness of carbon taxes relative to a fixed preference counter case. Additionally, the cultural multiplier is found to be stronger than the social multiplier, particularly when the carbon tax rate is low. This is due to the consensus-forming effect of the cultural multiplier, resulting in greater difficulty in individuals sustaining fringe consumption behaviours. However, this effect vanishes for stronger carbon tax rates where nuances between different social imitation schemes are dominated by the magnitude of the price signal.

Additionally, we show that with increasing substitutability of low- and high-carbon goods, the cultural multiplier strengthens, whilst the social multiplier weakens. In the absence of behavioural spillovers across different consumption categories, high substitutability amplifies polarization in consumption, reinforcing entrenched high-carbon preferences. Therefore, individuals with high-carbon preferences become less responsive to social influence due to their like-minded social bubble. Our findings suggest that greater similarity in environmental identity among peers connected via social networks increases the effective carbon tax rate required to reach emission reduction targets. Additionally, a richer or more diverse representation of lifestyles, achieved by increasing the number of consumption categories, enhances the strength of the cultural multiplier at low carbon tax levels. Our global sensitivity analysis confirms the robustness of our results over large parameter ranges.

In future research, the model could be extended to include rebound between consumption categories due to context-dependent preferences. This might include the possibility of low-carbon consumption in one consumption category leading to increased emissions in other areas due to moral licensing effects (Gholamzadehmir et al., 2019). Moreover, one could extend the model with a utility function possessing non-homothetic preferences to capture the heterogeneous behaviour of agents in different expenditure deciles. This would capture the ease with which more wealthy individuals can switch to low-carbon alternatives (Oswald, Owen, et al., 2020; Oswald, Millward-Hopkins, et al., 2023).

By fostering stronger pro-environmental identities, policymakers can leverage the cultural multiplier to reduce an effective carbon tax rate, contributing to greater policy support. This may be achieved through the introduction of complementary policies to a carbon tax. This may take the form of extending current visions of low-carbon lifestyles to be more systemic or rich in detail, including consumption in a high number of categories. For example information provision policies such as eco-labelling may correct

misinformation on the true carbon impact of less socially salient consumption categories. Alternatively, increasing the substitutability of low- and high-carbon alternatives both through technological improvements (e.g. widespread charging infrastructure and higher battery capacity to alleviate EV range anxiety (Pevec et al., 2020)) and nudge techniques to increase the social acceptability of low-carbon alternatives (e.g. plant-based meat substitutes (Edenbrandt and Lagerkvist, 2021; Coucke et al., 2022)). Additionally, policymakers should be mindful of network structures in which social imitation occurs when evaluating the expected effect of carbon taxation, as high similarity in pro-environmental identities amongst communities can act as roadblocks to decarbonization.

Code availability statement

The model code and documentation is available at:

<https://www.comses.net/codebases/b00ad1a3-dc5a-4610-83df-e869650f2714/releases/1.0.0/>

References

Albert, Réka and Albert-László Barabási (2002). “Statistical mechanics of complex networks”. In: *Reviews of modern physics* 74.1, p. 47.

Andersson, David and Jonas Nässén (2022). “Measuring the direct and indirect effects of low-carbon lifestyles using consumption data”. In: *Journal of Cleaner Production*, p. 135739.

Axelrod, Robert (1997). “The dissemination of culture: A model with local convergence and global polarization”. In: *Journal of conflict resolution* 41.2, pp. 203–226.

Bezin, Emeline (2019). “The economics of green consumption, cultural transmission and sustainable technological change”. In: *Journal of Economic Theory* 181, pp. 497–546.

Coucke, Nicky, Iris Vermeir, Hendrik Slabbinck, Maggie Geuens, and Ziad Choueiki (2022). “How to reduce agri-environmental impacts on ecosystem services: the role of nudging techniques to increase purchase of plant-based meat substitutes”. In: *Ecosystem Services* 56, p. 101444.

Dandekar, Pranav, Ashish Goel, and David T Lee (2013). “Biased assimilation, homophily, and the dynamics of polarization”. In: *Proceedings of the National Academy of Sciences* 110.15, pp. 5791–5796.

Davis, Taylor, Erin P Hennes, and Leigh Raymond (2018). “Cultural evolution of normative motivations for sustainable behaviour”. In: *Nature Sustainability* 1.5, pp. 218–224.

DeGroot, Morris H (1974). “Reaching a consensus”. In: *Journal of the American Statistical Association* 69.345, pp. 118–121.

Edenbrandt, Anna Kristina and Carl-Johan Lagerkvist (2021). “Is food labelling effective in reducing climate impact by encouraging the substitution of protein sources?” In: *Food Policy* 101, p. 102097.

Epstein, Joshua M and Robert Axtell (1996). *Growing artificial societies: social science from the bottom up*. Brookings Institution Press.

Falkenberg, Max, Alessandro Galeazzi, Maddalena Torricelli, Niccolò Di Marco, Francesca Larosa, Madalina Sas, Amin Mekacher, Warren Pearce, Fabiana Zollo, Walter Quattrociocchi, et al. (2022). “Growing polarization around climate change on social media”. In: *Nature Climate Change* 12.12, pp. 1114–1121.

Flüe, Lukas von and Sonja Vogt (2024). "Integrating social learning and network formation for social tipping towards a sustainable future". In: *Current Opinion in Psychology*, p. 101915.

Foramitti, Joël, Ivan Savin, and Jeroen van den Bergh (2024). "How carbon pricing affects multiple human needs: An agent-based model analysis". In: *Ecological Economics* 217, p. 108070.

Gholamzadehmir, Maedeh, Paul Sparks, and Tom Farsides (2019). "Moral licensing, moral cleansing and pro-environmental behaviour: The moderating role of pro-environmental attitudes". In: *Journal of Environmental Psychology* 65, p. 101334.

Girod, Bastien, Detlef Peter van Vuuren, and Edgar G Hertwich (2014). "Climate policy through changing consumption choices: options and obstacles for reducing greenhouse gas emissions". In: *Global Environmental Change* 25, pp. 5–15.

Glaeser, Edward L, Bruce I Sacerdote, and Jose A Scheinkman (2003). "The social multiplier". In: *Journal of the European Economic Association* 1.2-3, pp. 345–353.

Hartmann, Patrick, Aitor Marcos, and Jose M Barrutia (2023). "Carbon tax salience counteracts price effects through moral licensing". In: *Global Environmental Change* 78, p. 102635.

Herman, Jon and Will Usher (Jan. 2017). "SALib: An open-source Python library for Sensitivity Analysis". In: *The Journal of Open Source Software* 2.9.

Holland, Paul W, Kathryn Blackmond Laskey, and Samuel Leinhardt (1983). "Stochastic blockmodels: First steps". In: *Social networks* 5.2, pp. 109–137.

Kaaronen, Roope Oskari and Nikita Strelkovskii (2020). "Cultural evolution of sustainable behaviors: Pro-environmental tipping points in an agent-based model". In: *One Earth* 2.1, pp. 85–97.

Koessler, Ann-Kathrin, Stefanie Engel, et al. (2021). "Policies as information carriers: How environmental policies may change beliefs and consequent behavior". In: *International Review of Environmental and Resource Economics* 15.1-2, pp. 1–31.

Konc, Théo and Ivan Savin (2019). "Social reinforcement with weighted interactions". In: *Physical Review E* 100.2, p. 022305.

Konc, Théo, Ivan Savin, and Jeroen van den Bergh (2021). "The social multiplier of environmental policy: Application to carbon taxation". In: *Journal of Environmental Economics and Management* 105, p. 102396.

Kuperman, Marcelo N (2006). "Cultural propagation on social networks". In: *Physical Review E* 73.4, p. 046139.

Lanzini, Pietro and John Thøgersen (2014). "Behavioural spillover in the environmental domain: An intervention study". In: *Journal of Environmental Psychology* 40, pp. 381–390.

Mattauch, Linus, Cameron Hepburn, Fiona Spuler, and Nicholas Stern (2022). "The economics of climate change with endogenous preferences". In: *Resource and Energy Economics* 69, p. 101312.

Mattioli, Giulio, Milena Büchs, and Joachim Scheiner (2023). "Who flies but never drives? Highlighting diversity among high emitters for passenger transport in England". In: *Energy Research & Social Science* 99, p. 103057.

Nielsen, Kristian S, Kimberly A Nicholas, Felix Creutzig, Thomas Dietz, and Paul C Stern (2021). "The role of high-socioeconomic-status people in locking in or rapidly reducing energy-driven greenhouse gas emissions". In: *Nature Energy* 6.11, pp. 1011–1016.

Oswald, Yannick, Joel Millward-Hopkins, Julia K Steinberger, Anne Owen, and Diana Ivanova (2023). “Luxury-focused carbon taxation improves fairness of climate policy”. In: *One Earth*.

Oswald, Yannick, Anne Owen, and Julia K Steinberger (2020). “Large inequality in international and intranational energy footprints between income groups and across consumption categories”. In: *Nature Energy* 5.3, pp. 231–239.

Pevec, Dario, Jurica Babic, Arthur Carvalho, Yashar Ghiassi-Farrokhfal, Wolfgang Ketterer, and Vedran Podobnik (2020). “A survey-based assessment of how existing and potential electric vehicle owners perceive range anxiety”. In: *Journal of cleaner Production* 276, p. 122779.

Sato, Kazuo (1967). “A two-level constant-elasticity-of-substitution production function”. In: *The Review of Economic Studies* 34.2, pp. 201–218.

Schumacher, Ingmar (2015). “The endogenous formation of an environmental culture”. In: *European Economic Review* 76, pp. 200–221.

Sobol, Ilya M (2001). “Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates”. In: *Mathematics and computers in simulation* 55.1-3, pp. 271–280.

Sovacool, Benjamin K and Steve Griffiths (2020). “Culture and low-carbon energy transitions”. In: *Nature Sustainability* 3.9, pp. 685–693.

Torren-Peraire, Daniel, Ivan Savin, and Jeroen van den Bergh (2024). “An Agent-Based Model of Cultural Change for a Low-Carbon Transition”. In: *Journal of Artificial Societies and Social Simulation* 27.1, p. 13. ISSN: 1460-7425.

Truelove, Heather Barnes, Amanda R Carrico, Elke U Weber, Kaitlin Toner Raimi, and Michael P Vandenbergh (2014). “Positive and negative spillover of pro-environmental behavior: An integrative review and theoretical framework”. In: *Global Environmental Change* 29, pp. 127–138.

Van der Werff, Ellen, Linda Steg, and Kees Keizer (2014). “I am what I am, by looking past the present: the influence of biospheric values and past behavior on environmental self-identity”. In: *Environment and behavior* 46.5, pp. 626–657.

Wang, Tiantian, Bo Shen, Cecilia Han Springer, and Jing Hou (2021). “What prevents us from taking low-carbon actions? A comprehensive review of influencing factors affecting low-carbon behaviors”. In: *Energy Research & Social Science* 71, p. 101844.

Watts, Duncan J and Steven H Strogatz (1998). “Collective dynamics of ‘small-world’networks”. In: *Nature* 393.6684, pp. 440–442.

Westin, Kerstin, Johan Jansson, and Annika Nordlund (2018). “The importance of socio-demographic characteristics, geographic setting, and attitudes for adoption of electric vehicles in Sweden”. In: *Travel Behaviour and Society* 13, pp. 118–127.

Appendix

3.A Analytical results for the NCES utility function

In the M -dimensional case we have low and high-carbon goods for each consumption category, L_m and H_m , with an associated preference between the two A_m , and substitutability between goods σ . Between categories there is a further preference for consumption a_m where $\sum_{m=1}^M a_m = 1$ and substitutability across categories ν .

$$\max_{L_1, \dots, L_M, H_1, \dots, H_M} U(L_1, \dots, L_M, H_1, \dots, H_M, a_1, \dots, a_M, A_1, \dots, A_M, \sigma_1, \dots, \sigma, \nu) \quad (3.17)$$

the utility function to maximise is given by,

$$U = \left(\sum_{m=1}^M a_m U_m^\omega \right)^{\frac{1}{\omega}}, \quad (3.18)$$

where the psuedo-utility U_m is given by

$$U_m(L_m, H_m, A_m, \sigma) = (A_m L_m^\psi + (1 - A_m) H_m^\psi)^{\frac{1}{\psi}} \quad (3.19)$$

to simplify notation of the substitutabilities between low- and high-carbon goods for each category σ , and the between categories ν , we use use $\psi = \frac{\sigma-1}{\sigma}$ and $\omega = \frac{\nu-1}{\nu}$. This is subject to the budget constraint,

$$B = \sum_{m=1}^M L_m P_{L,m} + H_m P_{H,m} \quad (3.20)$$

To derive the demand functions for the utility function we require the Lagrangian for the system, given by

$$\mathcal{L} = \left(\sum_{m=1}^M a_m U_m^\omega \right)^{\frac{1}{\omega}} - \lambda \left(\sum_{m=1}^M L_m P_{L,m} + H_m P_{H,m} - B \right), \quad (3.21)$$

This produces general first-order conditions of low and high-carbon goods

$$\frac{\partial \mathcal{L}}{\partial L_m} = a_m \left(\sum_{m=1}^M a_m U_m^\omega \right)^{\frac{1}{\omega}-1} U_m^{\omega-1} \frac{\partial U_m}{\partial L_m} - \lambda P_{L,m} = 0 \quad (3.22)$$

$$\frac{\partial \mathcal{L}}{\partial H_m} = a_m \left(\sum_{m=1}^M a_m U_m^\omega \right)^{\frac{1}{\omega}-1} U_m^{\omega-1} \frac{\partial U_m}{\partial H_m} - \lambda P_{H,m} = 0 \quad (3.23)$$

In order to find L_m in terms of H_m we use the first order conditions with respect to the low and high-carbon good of the same category (same top level CES nest), re-arranging for λ and equating the two

$$\frac{1}{P_{L,m}} \frac{\partial U_m}{\partial L_m} = \frac{1}{P_{H,m}} \frac{\partial U_m}{\partial H_m} \quad (3.24)$$

$$\frac{\left(\frac{\partial U_m}{\partial H_m}\right)}{\left(\frac{\partial U_m}{\partial L_m}\right)} = \frac{P_{H,m}}{P_{L,m}} \quad (3.25)$$

We now produce the derivative of the psuedo-utilites U_m with respect to H_m and L_m

$$\frac{\partial U_m}{\partial L_m} = (A_m L_m^\psi + (1 - A_m) H_m^\psi)^{\frac{1}{\psi} - 1} A_m L_m^{\psi-1} \quad (3.26)$$

$$\frac{\partial U_m}{\partial H_m} = (A_m L_m^\psi + (1 - A_m) (H_m - h_m)^\psi)^{\frac{1}{\psi} - 1} (1 - A_m) H_m^{\psi-1} \quad (3.27)$$

Note that Equations 3.26 and 3.27 do not contain any between category terms. Substituting in the partial differentials of U_m with respect to H_m and L_m into our equated first order conditions in Equation 3.25 we produce a relationship between the quantity of low-carbon L_m and high-carbon good H_m

$$\frac{(1 - A_m) H_m^{\psi-1} (A_m L_m^\psi + (1 - A_m) H_m^\psi)^{\frac{1}{\psi} - 1}}{A_m L_m^{\psi-1} (A_m L_m^\psi + (1 - A_m) H_m^\psi)^{\frac{1}{\psi} - 1}} = \frac{P_{H,m}}{P_{L,m}} \quad (3.28)$$

$$\frac{H_m^{\psi-1}}{L_m^{\psi-1}} = \frac{P_{H,m} A_m}{P_{L,m} (1 - A_m)} \quad (3.29)$$

$$\frac{L_m}{H_m} = \left(\frac{P_{H,m} A_m}{P_{L,m} (1 - A_m)} \right)^{\frac{-1}{\psi-1}}, \quad (3.30)$$

in terms of substituabilities between low- and high-carbon goods, using the property $\sigma = \frac{-1}{\psi-1}$, the general ratio between the low and high-carbon good for the m^{th} category is defined as

$$\Omega_m = \frac{L_m}{H_m} = \left(\frac{P_{H,m} A_m}{P_{L,m} (1 - A_m)} \right)^\sigma. \quad (3.31)$$

Next, we compare low-carbon goods from different categories to derive the ratio between high-carbon goods for two different categories (H_p and H_q). Re-arranging the first-order conditions of two low-carbon goods and equating them, where p, q are dummy variables for the m^{th} category,

$$\frac{1}{P_{L,p}} \frac{\partial U}{\partial L_p} = \frac{1}{P_{L,q}} \frac{\partial U}{\partial L_q} \quad (3.32)$$

$$\frac{1}{P_{L,p}} a_p \left(\sum_{m=1}^M a_m U_m^\omega \right)^{\frac{1}{\omega} - 1} U_p^{\omega-1} \frac{\partial U_p}{\partial L_p} = \frac{1}{P_{L,q}} a_q \left(\sum_{m=1}^M a_m U_m^\omega \right)^{\frac{1}{\omega} - 1} U_q^{\omega-1} \frac{\partial U_q}{\partial L_q} \quad (3.33)$$

$$\frac{1}{P_{L,p}} a_p U_p^{\omega-1} \frac{\partial U_p}{\partial L_p} = \frac{1}{P_{L,q}} a_q U_q^{\omega-1} \frac{\partial U_q}{\partial L_q} \quad (3.34)$$

Substituting our expression of the partial differential from Equation 3.26 and the psuedo-utility U_m from Equation 3.19,

$$\frac{1}{P_{L,p}} a_p (A_p L_p^{\psi_p} + (1 - A_p) H_p^{\psi_p})^{\frac{\omega-1}{\psi_p}} \frac{\partial U_p}{\partial L_p} = \frac{1}{P_{L,q}} a_q (A_q L_q^{\psi_q} + (1 - A_q) H_q^{\psi_q})^{\frac{\omega-1}{\psi_q}} \frac{\partial U_q}{\partial L_q} \quad (3.35)$$

$$\begin{aligned}
& \frac{1}{P_{L,p}} a_p (A_p L_p^{\psi_p} + (1 - A_p) H_p^{\psi_p})^{\frac{\omega-1}{\psi_p}} (A_p L_p^{\psi_p} + (1 - A_p) H_p^{\psi_p})^{\frac{1}{\psi_p}-1} A_p L_p^{\psi_p-1} \\
&= \frac{1}{P_{L,q}} a_q (A_q L_q^{\psi_q} + (1 - A_q) H_q^{\psi_q})^{\frac{\omega-1}{\psi_q}} (A_q L_q^{\psi_q} + (1 - A_q) H_q^{\psi_q})^{\frac{1}{\psi_q}-1} A_q L_q^{\psi_q-1}
\end{aligned} \tag{3.36}$$

$$\begin{aligned}
& \frac{1}{P_{L,p}} a_p A_p L_p^{\psi_p-1} (A_p L_p^{\psi_p} + (1 - A_p) H_p^{\psi_p})^{\frac{\omega-1}{\psi_p} + \frac{1}{\psi_p}-1} \\
&= \frac{1}{P_{L,q}} a_q A_q L_q^{\psi_q-1} (A_q L_q^{\psi_q} + (1 - A_q) H_q^{\psi_q})^{\frac{\omega-1}{\psi_q} + \frac{1}{\psi_q}-1}
\end{aligned} \tag{3.37}$$

$$\frac{1}{P_{L,p}} a_p A_p L_p^{\psi_p-1} (A_p L_p^{\psi_p} + (1 - A_p) H_p^{\psi_p})^{\frac{\omega}{\psi_p}-1} = \frac{1}{P_{L,q}} a_q A_q L_q^{\psi_q-1} (A_q L_q^{\psi_q} + (1 - A_q) H_q^{\psi_q})^{\frac{\omega}{\psi_q}-1} \tag{3.38}$$

Substituting in the general ratio of low to high-carbon goods within a branch, $L_m = H_m \Omega_m$,

$$\begin{aligned}
& \frac{1}{P_{L,p}} a_p A_p \Omega_p^{\psi_p-1} H_p^{\psi_p-1} (A_p \Omega_p^{\psi_p} H_p^{\psi_p} + (1 - A_p) H_p^{\psi_p})^{\frac{\omega}{\psi_p}-1} = \frac{1}{P_{L,q}} a_q A_q \Omega_q^{\psi_q-1} H_q^{\psi_q-1} (A_q \Omega_q^{\psi_q} H_q^{\psi_q} + (1 - A_q) H_q^{\psi_q})^{\frac{\omega}{\psi_q}-1} \\
& \frac{1}{P_{L,p}} a_p A_p \Omega_p^{\psi_p-1} H_p^{\omega-1} (A_p \Omega_p^{\psi_p} + 1 - A_p)^{\frac{\omega}{\psi_p}-1} = \frac{1}{P_{L,q}} a_q A_q \Omega_q^{\psi_q-1} H_q^{\omega-1} (A_q \Omega_q^{\psi_q} + 1 - A_q)^{\frac{\omega}{\psi_q}-1}
\end{aligned} \tag{3.39}$$

Now gathering high-carbon consumption terms,

$$\frac{H_p^{\omega-1}}{H_q^{\omega-1}} = \frac{a_q A_q \Omega_q^{\psi_q-1} (A_q \Omega_q^{\psi_q} + 1 - A_q)^{\frac{\omega}{\psi_q}-1}}{P_{L,q}} \frac{P_{L,p}}{a_p A_p \Omega_p^{\psi_p-1} (A_p \Omega_p^{\psi_p} + 1 - A_p)^{\frac{\omega}{\psi_p}-1}} \tag{3.41}$$

$$\frac{H_p}{H_q} = \left(\frac{a_q A_q \Omega_q^{\psi_q-1} (A_q \Omega_q^{\psi_q} + 1 - A_q)^{\frac{\omega}{\psi_q}-1}}{P_{L,q}} \frac{P_{L,p}}{a_p A_p \Omega_p^{\psi_p-1} (A_p \Omega_p^{\psi_p} + 1 - A_p)^{\frac{\omega}{\psi_p}-1}} \right)^{\frac{1}{\omega-1}} \tag{3.42}$$

In terms of category substitutabilities the ratio between the high-carbon quantities from different categories is given by, where we substitute in $\psi = \frac{\sigma-1}{\sigma}$ and $\omega = \frac{\nu-1}{\nu}$,

$$\frac{H_p}{H_q} = \left(\frac{P_{L,q} \Omega_q^{\frac{1}{\sigma_q}}}{a_q A_q \left(A_q \Omega_q^{\frac{\sigma_q-1}{\sigma_q}} + 1 - A_q \right)^{\frac{\nu-\sigma_q}{\nu(\sigma_q-1)}}} \frac{a_p A_p \left(A_p \Omega_p^{\psi_p} + 1 - A_p \right)^{\frac{\nu-\sigma_p}{\nu(\sigma_p-1)}}}{P_{L,p} \Omega_p^{\frac{\sigma_p-1}{\sigma_p}}} \right)^\nu \tag{3.43}$$

To simplify notation of the quantities we introduce an interaction term χ_m

$$\chi_m = \left(\frac{a_m A_m}{P_{L,m} \Omega_m^{\frac{1}{\sigma}}} \right)^\nu \left(A_m \Omega_m^{\frac{\sigma-1}{\sigma}} + 1 - A_m \right)^{\frac{\nu-\sigma}{(\sigma-1)}} \tag{3.44}$$

such that the quantity of dummy category p may be expressed in terms of the quantity of dummy category q

$$H_p = \left(\frac{\chi_p}{\chi_q} \right) H_q \tag{3.45}$$

Lastly to derive demand functions for the low- and high-carbon goods in terms of preferences, prices and substitutabilities we consider the budget constraint and use previous definition of low to high-carbon goods from Equation 3.31

$$B = \sum_{p=1}^M L_p P_{L,p} + H_p P_{H,p} \quad (3.46)$$

$$= \sum_{p=1}^M (H_p \Omega_p P_{L,p} + H_p P_{H,p}) \quad (3.47)$$

$$= \sum_{p=1}^M H_p (\Omega_p P_{L,p} + P_{H,p}) \quad (3.48)$$

Note that again p is a dummy variable representing any category. Now substituting in the interaction term χ_m between different categories defined in Equation 3.44 we express the high-carbon quantity of a given dummy category q in terms of the preferences, prices and substitutabilities,

$$B = \sum_{p=1}^M \left(\frac{\chi_p}{\chi_q} H_q \right) (\Omega_p P_{L,p} + P_{H,p}) \quad (3.49)$$

$$= \sum_{p=1}^M \frac{\chi_p}{\chi_q} (H_q (\Omega_p P_{L,p} + P_{H,p})) \quad (3.50)$$

$$= \frac{H_q}{\chi_q} \sum_{p=1}^M \chi_p (\Omega_p P_{L,p} + P_{H,p}) \quad (3.51)$$

$$H_q = \frac{\chi_q B}{\sum_{p=1}^M \chi_p (\Omega_p P_{L,p} + P_{H,p})} \quad (3.52)$$

Thus the quantity of the m^{th} good is given by,

$$H_m = \frac{B \chi_m}{Z} \quad (3.53)$$

$$L_m = \frac{B \Omega_m \chi_m}{Z} \quad (3.54)$$

where to simplify notation Z is defined as,

$$Z = \sum_{m=1}^M \chi_m (\Omega_m P_{L,m} + P_{H,m}) \quad (3.55)$$

this serves as a normalization term across categories.

3.B Imitation of consumption

The existence of a social multiplier relies on two key features regarding the model of social imitation; firstly, that the preferences of individuals for low-carbon goods are not observable and secondly that the utility function is not common knowledge. It is important

to note that in the case that either of these assumptions does not hold the social multiplier effect vanishes, as the carbon tax no longer has a channel through which to affect preferences. However, if we assume that some portion of social information is, in fact, a direct observation of preferences then we still find a non-linearity in decarbonisation if that signal is not entirely preference-based information.

To understand how the quantities of low- and high-carbon goods consumed affect the social imitation process, we require an understanding of how changes in preferences result in changes to low-carbon consumption ratio $C_{t,i,m}$ as a function of prices and substitutability.

$$C_{t,i,m} = \frac{L_{t,i,m}}{L_{t,j,m} + H_{t,i,m}} \quad (3.56)$$

$$= \frac{H_{t,i,m} \Omega_{t,i,m}}{H_{t,i,m} \Omega_{t,i,m} + H_{t,i,m}} \quad (3.57)$$

$$= \frac{\Omega_{t,i,m}}{\Omega_{t,i,m} + 1} \quad (3.58)$$

Now substituting in the ratio of low to high-carbon consumption $\Omega_{t,i,m}$

$$C_{t,i,m} = \frac{\left(\frac{P_{H,m} A_{t,i,m}}{P_{L,m}(1-A_{t,i,m})}\right)^\sigma}{\left(\frac{P_{H,m} A_{t,i,m}}{P_{L,m}(1-A_{t,i,m})}\right)^\sigma + 1} \quad (3.59)$$

To simplify notation the ratio of prices between low- and high-carbon goods \bar{P}_m is defined as

$$\bar{P}_m = \frac{P_{L,m}}{P_{H,m}} \quad (3.60)$$

We substitute in the price ratio to obtain a simplified low-carbon consumption proportion

$$C_{t,i,m} = \frac{\left(\frac{A_{t,i,m}}{\bar{P}_m(1-A_{t,i,m})}\right)^\sigma}{\left(\frac{A_{t,i,m}}{\bar{P}_m(1-A_{t,i,m})}\right)^\sigma + 1} \quad (3.61)$$

$$= \frac{A_{t,i,m}^\sigma}{A_{t,i,m}^\sigma + (\bar{P}_m(1 - A_{t,i,m}))^\sigma} \quad (3.62)$$

The smaller the value of the price ratio \bar{P}_m the smaller the value of the low-carbon preference $A_{t,i,m}$ required to induce a complete switch to low-carbon good consumption in that category. Figure 3.B.1 shows the dependence of $C_{t,i,m}$ on $A_{t,i,m}$. Small differences in preferences for goods can lead to large changes in consumption proportions due to the nonlinear impact of substitutability between the goods and price differences.




Figure 3.B.1: The proportion of total consumption assigned to low-carbon good within a category $C_{t,i,m}$ as a function of the preference for that low-carbon good $A_{t,i,m}$ and price ratio between low- and high-carbon goods \bar{P}_m .

Additionally, under the conditions $\sigma \rightarrow 1$ and prices between low- and high-carbon goods equal $\bar{P}_m = 1$ then

$$C_{t,i,m} = A_{t,i,m} \quad (3.63)$$

Therefore, under these conditions the preferences dynamics collapses to those studied in Torren-Peraire et al. (2024). On the other hand, when goods are perfect substitutes $\sigma \rightarrow \infty$ then $C_{t,i,m}$ tends to a step function in terms of $A_{t,i,m}$, where the location of the step in preference space is given by

$$A_{step,m} = \frac{\bar{P}_m - 1}{\bar{P}_m}. \quad (3.64)$$

3.C Additional simulation results




Figure 3.C.1: Environmental identity time-series at two carbon taxes for the small-world network.




Figure 3.C.2: *Cumulative emissions for the case of fixed preferences (blue), cultural multiplier (green) and social multiplier (orange) using a Stochastic Block Model and scale free network. Shaded region indicated 95% confidence interval.*




Figure 3.C.3: *Carbon tax reduction relative to fixed preferences for cultural multiplier (green) and social multiplier (orange) using a Stochastic Block Model and scale free network. Shaded region indicated 95% confidence interval for 100 stochastic runs.*




Figure 3.C.4: *Varying substitutability between the low- and high-carbon good alternatives.*




Figure 3.C.5: *First order Sobol sensitivity analysis for 128 parameter values for a total of 184,320 experiments*




Figure 3.C.6: *Cumulative emissions for increasing number of consumption categories at different carbon tax values, for a total of 2150 runs.*

3.C.1 Network structure and homophily in initial environmental identity

To investigate the role of initial homophily (similarity in environmental identity) in shaping decarbonisation dynamics, we use the SW network and SBM. Specifically, we examine whether initial homophily creates barriers or promotes social tipping points towards the adoption of low-carbon alternatives. On the other hand, for the SF network, the large asymmetries in a number of social connections mean that similarities in the initial environmental identity of central agents are of key interest. We label this concentration of similar environmental identities in high-degree nodes as a low- or high-carbon hegemony. Note that this section builds on the social learning module whereby individuals socially imitate in a weighted fashion based on similarity in identity (cultural multiplier). Therefore, we consider both how initial and dynamic homophily in environmental identity change the effectiveness of a carbon tax simultaneously.

The results in Figure 3.C.7 show that the introduction of homophily reduces the strength of the cultural multiplier both in the SW and SBM networks. Greater initial homophily in environmental identity sustains high-carbon consumption practices (relative to the no or low homophily counter-case) of individuals as they closely imitate each other. For the SBM, greater homophily means that low values of the carbon tax are unable to induce a mass change in consumption behaviours. However, when a critical carbon tax rate is reached, the system tips towards low-carbon consumption.

This tipping behaviour is a result of the block structure of the network, which allows for a mixed distribution of preferences both between and within block communities. One community can exhibit high-carbon consumption while the other adopts low-carbon behaviours, or each community can have a mix of both high- and low-carbon consumption. Due to this effect, the cultural multiplier is negative for low tax values as greater decarbonisation would be achieved with fixed preferences (dashed black line). In contrast, for the SW network, we see a more gradual decline in emissions due to a more homogenous distribution of the node degree in the network, meaning no single individual or small community can tip the system towards low-carbon consumption. This is also seen in the lower magnitude of price elasticities for the SW network in Appendix Figure 3.C.9, relative to the SBM.

In some social contexts, such as online social networks, peer influence can be highly asymmetric. To capture this, we introduce the concept of hegemony, where a high-carbon hegemony reflects a concentration of high-carbon environmental identities amongst the most connected individuals. Figure 3.C.8 shows how, in the case of low-carbon hegemony, even without a carbon tax, emissions are significantly reduced relative to the no-homophily case. This occurs because centrally placed individuals exert strong influence over many neighbours with preferences for high-carbon consumption while selectively imitating those with similar environmental identities, thereby minimising exposure to opposing behaviours. Under high-carbon hegemony, the cultural multiplier is negative for low-carbon tax values, impeding decarbonisation efforts, but may become positive for a sufficiently strong carbon tax signal. In the case of high-carbon hegemony, the social network shifts towards low-carbon consumption from the periphery to the highly connected center. Due to this the SF network has a greater price elasticity than the SW network but lower than the SBM, see Appendix Figure 3.C.9.




Figure 3.C.7: *Emissions reduction due to carbon taxation for different degrees of initial homophily in environmental identity, for a total of 18,000 experiments. The dashed line is the mean emissions of 100 stochastic runs for the case of fixed preferences.*




Figure 3.C.8: *Cumulative emissions for a scale-free network where central nodes are grouped by environmental identity at the start of the experiment, for a total of 9,000 experiments. The dashed line is the mean emissions of 100 stochastic runs for the case of fixed preferences.*




Figure 3.C.9: *Price elasticities for three network structures with different distributions of initial environmental identity (homophily and hegemony).*

Chapter 4

Driving in the wrong direction? A co-evolutionary model of EV adoption and innovation*

4.1 Introduction

In 2019, 10% of global greenhouse gas emissions were attributed to road transport, which represented 70% of all carbon emissions related to transport. Therefore, the environmental sustainability of future generations requires large scale reductions in the carbon footprint of road transport. The large-scale adoption of electric vehicles (EVs), to replace internal combustion engine vehicles (ICEVs), is an indispensable condition for the transition to sustainable mobility in car-dependent urbanizations (Jaramillo et al., 2022, p 1078). This requires policy incentives to be cost-effective, environmentally sustainable and accepted by the public. However, the co-evolving socioeconomic interactions between consumer adoption of new products and producer investments in innovations make difficult for policy makers to anticipate the outcomes of the policies before implementing them. This is relevant not only for designing policies that are effective in promoting electric vehicle adoption, but also to avoid unintended negative consequences on other variables (Herrmann and Savin, 2017; Nuñez-Jimenez et al., 2022).

Agent-based simulation models (ABMs) are an appealing method to address public policy issues due to their bottom-up structure that accounts for agent heterogeneity and the complex interplay of individual behaviour. These models have offered useful policy design laboratories for stimulating the diffusion of EVs (Domarchi and Cherchi, 2023; Mehdizadeh et al., 2022), often calibrated for the data of the United States (Brown, 2013; Noori and Tatari, 2016; Adepetu et al., 2016; Sun et al., 2019; Spangher et al., 2019; Ledna et al., 2022). In the context of mobility transitions, the study of social networks and opinion formation dynamics has increased in popularity (McCoy and Lyons, 2014; Silvia and Krause, 2016; Li et al., 2019; Feng et al., 2019; He et al., 2020; Ning et al., 2020; Lee and Brown, 2021; Zhang et al., 2022). The literature of ABMs for EV diffusion also provides studies addressing the role of directed innovation in promoting more sustainable modes of transportation (Windrum et al., 2009; Greene et al., 2014; Sun et al., 2019; Fan and Li, 2025).

*The chapter was co-authored with Miquel Bassart-i-Lore. I coded the model, whilst Miquel led the manuscript writing. Model design, analysis and editing was conducted jointly.

Despite the growing literature, the interaction between social dynamics and innovation decisions remains underexplored. Specifically, the interplay between consumer adoption and firm innovation presents important considerations for policy analysis due to their co-evolving nature. Firms direct their innovation efforts to electric vehicles based on the size of its customer base. In turn, people are willing to adopt EVs if their attributes and price can compete with ICEVs. Without policy intervention, this interaction slows down the deployment of new products. However, policies motivating the early adoption and investment of EVs can trigger a virtuous circle that accelerates the transition. In order to be sustainable, such transition requires minimising policy costs, welfare losses, and greenhouse gas rebound effects. Therefore, our paper addresses the following research question: What policy combinations that meet adoption targets can best balance the trade-offs between costs, consumer utility, and emissions?

This paper provides a framework to support policy design accounting for the co-evolution of adoption and innovation. We introduce an ABM combining a multinomial logit model for consumer choices (Østli et al., 2017), a social imitation model for the diffusion of electric vehicles (Rogers, 2003), and an *NK* model for the directed innovation efforts of manufacturers (Windrum et al., 2009). These elements allow capturing the interaction between social imitation among heterogeneous car users and profit-oriented investment choices. Furthermore, we include a stylized representation of the market for used cars given its paramount importance in the US car market.

To calibrate the model, we set parameters to replicate the conditions of the state of California from 2001 to 2023, due to its rich data availability on mobility. Furthermore, California is an example of a car-dependent territory committed to a reduction in carbon emissions as part of its Zero-Emission Vehicle (ZEV) requirement, adopted by the California Air Resources Board. This mandate is complemented by Federal and State rebates supporting the adoption of EVs.

The model is used to test policy scenarios for the period 2024-2035. Particularly, we test a carbon price, a rebate for new and used EVs, an electricity subsidy and an EV production subsidy. Policies are evaluated in their ability to reach a target adoption rate and compared on their impact on total net policy costs, cumulative carbon emissions, and consumer utility. In this way, we study the induced trade-offs associated with each policy instrument. Moreover by combining the instruments in pairs we study how their performance can be improved. Finally, we show the stability of the EV transition by simulating for the period 2036-2050 after removing policy intervention.

The remainder of the paper is organized as follows. In Section 4.2 we formulate the automotive sector model of endogenous innovation and social imitation in the context of climate policy. Section 4.3 provides the details of our calibration strategy to externally validate the model outputs with respect to Californian data. In Section 4.4, we analyze the results of the model. Finally, Section 4.5 concludes and outlines future research avenues.

4.2 The model

4.2.1 Overview

The present model studies policies for the diffusion of EVs in the context where innovation and social imitation are co-evolving processes. The model features a market for new and used cars that are characterized by a car type (ICEV or EV), production cost, fuel

efficiency, battery or fuel tank size, and an abstract quality measure. The model is divided into four submodules: the discrete choice consumption submodule, the social imitation submodule, the innovation submodule and the used car market submodule. Figure 4.2.1 provides an overview of the interactions between each submodule. The model is in discrete time, where each time step represents a month. For model equations, see Appendix 4.A.

For the representation of the users' choice of car, we use a discrete choice logit model (McFadden, 1974), as in previous studies addressing the adoption of EVs (Eggers and Eggers, 2011; Shafiei et al., 2012; Østli et al., 2017). Heterogeneous car users meet their mobility needs either keeping their current car, purchasing a new or a used car. The evaluation of a car is based on its attributes and consumer preferences. We add to this literature by modelling car users that are forward-looking in estimating the lifecycle costs and emissions when evaluating a car.

We use the diffusion of innovations theory by Rogers (Rogers, 2003) to conceptualize the willingness to consider EVs. Users become willing to consider buying an EV after a sufficient number of related peers are using it. The number of neighbours threshold varies amongst car users, with its value corresponding to a spectrum between "innovators" and "laggards".

Car manufacturers select the car designs to put on sale based on their profit expectations throughout the lifecycle of the technology, evaluated at the prices consistent with profit maximization. New ICEV or EV vintages are discovered as a result of innovation. Following similar studies (Windrum et al., 2009), we employ the *NK* model (Kauffman, 1993) to represent the distribution of possible technologies. In this setting, the available car designs depend on previous *R&D* choices, leading to specialization patterns and emergent heterogeneity.

For the representation of the market for used cars, we propose a stylized model where used car dealers are consolidated into a single entity. The price of used cars follows the price of their respective most similar new car and applies a price discount proportional to the car age. Whenever a consumer buys a car, the replaced car is sold to the used car dealer at its market price with a discount. Used car dealers discard cars whose market price falls below its scraping value.

4.2.2 Car users

Car users drive a fixed and heterogeneous distance each month. Stochastically, users are sequentially activated to replace their car for one of new or used cars available on the market. Following a social imitation model where car users are placed in a small-world network, the willingness to consider an EV is updated at the end of the period. The model assumes one car for each driver, as the US car-to-driver ratio has remained stable at 1.15 from 2012 to 2023 (Bureau of Transportation Statistics).

Car choice

Car users compare the lifecycle utility of their own car with that of the available used and new cars the user is willing to consider. The utility of a car is given by the willingness to pay for its quality and range, each of them with diminishing returns, minus its lifecycle costs and emissions (Jaramillo et al., 2022, p.4). The lifecycle costs are given by the sum of the purchase price, net of the sale of the user's current car, and the discounted sum of the expected cost per kilometre. The lifecycle emissions are given by the manufacturing




Figure 4.2.1: *Model structure*

emissions if the car is new and the discounted sum of the expected emissions per kilometre. When estimating future costs and emissions related to driving, car users factor in the depreciation of fuel efficiency taking place at the end of each time-step at a fixed rate. Given the uncertainty in gasoline and electricity markets, car users assume the future prices and emissions of gasoline and electricity to be the same as in the present period.

Given a utility for each of the alternative car choices (including their own) car users select one using a logit model (McFadden, 1974) where the probability of selecting a car is proportional to its assigned utility.

Social imitation

To represent the influence of peers on the choice of car type, we build on the threshold models of EV adoption outlined in Eppstein et al., 2011; McCoy and Lyons, 2014; Wang et al., 2023. We consider a social network in which car users change their car purchasing behaviour based on the proportion of EV owned by their neighbours. This threshold is heterogeneous amongst car users, with a low threshold representing "innovators" who are willing to adopt new technologies, whilst a high threshold corresponds to "laggards" who will only consider moving away from the established technology once it is very common in their network neighbourhood.

The result of the threshold model is that car users may not adopt EVs even if those have high utility. This structure captures the gap between behaviour and attitudes or beliefs in which car users may have sufficient desire (high utility) to pursue pro-environmental behaviours but do not do so due to other limiting factors (Kollmuss and Agyeman, 2002), such as the lack of an established social norm. Thus, it is possible for car users to have a high willingness to pay for lower emissions and high emissions due to

driving an ICEV. On the other hand, with policy intervention, social imitation dynamics can amplify the policy effect by increasing the visibility of EVs (Konc et al., 2021).

Social interaction is structured using a small world network (Watts and Strogatz, 1998) to replicate word-of-mouth social interactions in highly clustered social circles where neighbours are also likely to become friends. This network is formed by first arranging car users in a ring, with each person connected to a homogenous number of nearest neighbours, a fixed number of connections each to their nearest neighbours. Then the symmetry of the network is broken up by introducing long-range weak ties across the network through probabilistic re-wiring of connections.

4.2.3 Car supply

Manufacturers sell multiple car designs that are produced on demand. Stochastically, they update the car types on sale and their respective prices in line with profit maximization. For simplicity, we assume car manufacturers face constant production costs and no capacity or financial constraints. With these assumptions, we ignore some relevant aspects for the transition, such as carbon lock-in (Unruh, 2000; Unruh, 2002) or lack of financial appetite from venture capitalists (Mazzucato, 2013; Mazzucato and Semieniuk, 2017). We choose this strong assumption based on the fact that the car industry is formed by firms that are integrated in global supply chains and financial markets, while our model is restricted to the local market.

Following the Segmentation, Targeting and Positioning framework (Kotler, 1989), manufacturers supply car designs targeting a specific consumer segment. A consumer segment is defined as a relatively homogeneous group of customers with particular needs and preferences (Lynn, 2012). In this model, segments are formed on the basis of EV acceptance, willingness to pay for quality and willingness to pay for emissions reduction. These are relevant determinants of car choice after controlling for socio-demographic attributes (Axsen and Kurani, 2013; Keyes and Crawford-Brown, 2018; Jang and Choi, 2021).

Technology choice and pricing

Manufacturers evaluate the expected profit at each consumer segment of each car type using known technologies. The expected profit is given by the marginal profit, the segment population, and the expected market share within the segment. This expected in-segment market share depends on the attributes of the car and the state of competition in that segment (over the last 12 months). Crucially, it is also consistent with the discrete choice consumption model, where the price is defined to maximize expected profits.

The choice of car types being supplied is resolved heuristically. Manufacturers include in their supply line the most profitable car targeted at a specific segment. Once the segment is satisfied, the selected price is locked. The remaining segments are satisfied by sequentially selecting the most profitable car, where the expected profits of unselected cars targeted at their most profitable remaining segment are compared to that of the selected cars with the profit-maximizing price for the selected segments. As a result, a car may be produced to satisfy multiple segments at a single price if it outperforms other designs.

By assuming that all manufacturers target all segments, we abstract from the brand-reputation considerations. This simplification allows us to focus on the interaction be-

tween innovation and adoption, especially in the context of transition policies in the car industry. However, to encourage specialization into specific segments, we limit the number of cars that may be produced by a single manufacturer to below the total number of segments and limit the number of technologies a manufacturer can retain in its memory.

Innovation

Stochastically, car manufacturers innovate and learn technology to produce a new car type with unique characteristics. We use *NK* models to conceptualize innovation (Levinthal, 1997; Rivkin, 2000; Baumann and Siggelkow, 2013), see Appendix 4.E. In this framework, technologies are distributed on a map where local movements lead to incremental improvements in product attributes. *NK* models allow the representation of complex technological landscapes with multiple local solutions, accounting for the path-dependent nature of innovation. This assumption allows the capture of heterogeneous abilities to adopt technologies, or absorptive capacity (Cohen and Levinthal, 1990), which are gained with path-dependent experience (Nelson and Winter, 1977).

Manufacturers are engaged in a stochastic parallel search in the landscapes for ICEVs and EVs. We limit innovation to the development of one single car, thus making the choice to develop ICEVs or EVs exclusive. This choice is made based on the profit expectation of using the candidate car, whose attributes are known to the manufacturer, priced at its most profitable segment. Mirroring consumer car choice, firms use a logit model to choose which innovative design to pursue, with an intensity of choice parameter λ that determines the balance of exploration and exploitation (Puram et al., 2015).

As in previous studies (Jain and Kogut, 2014), manufacturers discard previously searched solutions when forming the set of candidate technologies to explore. This set contains all the unexplored neighbouring technologies with respect to the past search location on each landscape. We include the past search location, allowing the possibility of not searching if there are no attractive alternatives. Once a new design has been researched and the memory capacity is exceeded, the oldest car design not in production is forgotten. The limited memory capacity is a cognitive limitation whereby the searching agent economizes cognitive power by forgetting unused solutions.

Used cars

The market for used cars represents the bulk of car purchases in the US¹, providing an affordable alternative to low-income users. This model chooses a parsimonious approach for the representation of the market for used cars. Used car dealers are consolidated into a single entity that sets the purchase and sale prices of used cars.

The sale price of a used car is a fraction of the price of its most similar car on the market for new cars. The fraction is inversely proportional to the used car age, where the price depreciation rate of EVs is larger than that of ICEVs (Schloter, 2022). The used car purchase price is a fraction of its sale price, given by a fixed target mark-up that serves as a proxy for the degree of monopsony in the market for used cars. Sale and purchase prices have an exogenous lower bound representing the car scrap value. Cars with prices lower than their scrap value are removed from circulation. To reduce computational cost, we impose a maximum capacity on the second-hand market after which the cars with the lowest value are removed.

¹Statista

4.3 Calibration and validation

We calibrate the model to the data for the state of California 2001 to 2023, where each model time step represents a month. For summary of values used see Appendix 4.B and for more details on data sources and justification see Appendix 4.C. All monetary units are expressed in 2020 US dollars. The price and emissions per kilowatt hour of gasoline and electricity during the validation period reflect real world monthly data, while the distribution of monthly distance driven is taken from survey data. In the case of distance travelled, we assume no rebound due to a change in car type (Zhang et al., 2025). We extrapolate secondary data from previous studies to calibrate the willingness to pay for range increases and emissions reductions, the boundaries of the car attributes in the NK landscape (except for quality), the car scrap value, the discount rate, and the price depreciation rate of used cars.

For other non-observable parameters, we follow past history-friendly models (Malerba et al., 2008; Herrmann and Savin, 2017) in avoiding excessive calibration. Instead, we fix the least sensitive parameters to some value at a reasonable order of magnitude and perform the respective sensitivity analyses afterwards. This applies to the parameter governing technological complexity, the research intensity of choice, the mark-up applied in the used car market, and the diminishing returns to utility with respect to quality. To capture the nuances between EV and ICE car production, we assume the battery size is correlated to unit production costs by a coefficient of .5, such that increases in EV range may come at a greater cost.

The remaining parameters are indirectly calibrated to replicate the trend of EV uptake and percentage of EV sales in California (see California Energy Commission data). Additionally, our indirect calibration keeps output variables such as prices, market concentration or average car age within the observed ranges described in Table 4.3.1. In Figure 4.H.1, we show a representative distribution of the willingness-to-pay parameters, distances and innovativeness values.

Table 4.3.1: Target variable ranges

Variable	Range	Source
Price for new cars	25th %: \$32,359.41, 75th %:\$57,784.66	(Grieco et al., 2024, Figure II)
Herfindahl-Hirshman index	(.11-.18)	(Grieco et al., 2024, Figure II)
Average car age between	(10-12) years	Bureau of Transportation Statistics

Based on the conducted Sobol sensitivity analysis, see Appendix 4.D, the key parameters affecting our target variables are the distributions of the willingness to pay for quality and of the imitation threshold, the consumer intensity of choice, and the fuel efficiency depreciation rate. Since the willingness to pay for the quality parameter mostly affects the price level, we determine its median value analytically in such a way that the optimal price of an ICEV with average features targeted at a consumer with average willingness

to pay equals the observed average price. Then, we distribute this value according to the Californian income distribution, assigning a higher willingness to pay to richer agents.

For the users' intensity of choice and the depreciation rate of car efficiency, we perform a grid search to match the stylized facts of car price range and age, iteratively narrowing the search space until matching the trends for EV uptake while keeping the other target variables within their observed ranges. Lastly, a beta distribution is used to define the heterogeneous values of car users' EV consideration thresholds. Given the high computation cost of the ABM, simulation-based Bayesian inference (Papamakarios and Murray, 2016; Cranmer et al., 2020; Dyer et al., 2024) is used to calibrate the beta distribution values to match EV uptake during the validation period.

We simulate a population of 3,000 car users and 16 car manufacturers (Grieco et al., 2024). The model simulates a burn-in period of 15 years (180 time-steps) before the validation period. The model is initialized with car manufacturers being placed at one Hamming distance from the worst technology when evaluated at the median segment, and car users owning a random car among those manufacturers can produce. During the burn-in period, manufacturers improve and sell ICEVs, while EV search and production is not permitted until the validation period.

In order to account for existing EV support policies, we feature a simplified version of the federal and state rebate assigned to EVs (California Air Resources Board, 1990a; California Air Resources Board, 1990c; California Air Resources Board, 1990b) from 2010-23, discounting a flat amount of \$10,000 to new EVs and \$1,000 to used EVs. In addition, the implemented carbon taxes are already included in the imported prices for gasoline and electricity. However, for the sake of simplicity, we do not model the existing fuel efficiency standards or the EV sale mandates stipulated in the ZEV Californian plan.

Figure 4.3.1 shows the model projections for the validation on EV uptake and percentage of EV sales (top left), prices of new and used EVs and ICEVs (top right), market concentration (bottom right), and average car age (bottom left). The simulated values are the average over 64 Monte Carlo simulations. These plots validate our indirect calibration exercise, given that the projections of EV uptake and sales closely match the real data and that prices, market concentration, and average car age are mostly within the ranges described in Table 4.3.1.

4.4 Policy experiments

4.4.1 Experimental setup

The objective of this research is to analyze policy programs aimed at achieving full electrification of cars by 2035. In this study, we focus exclusively on taxes and subsidies. While other policy instruments, such as regulations, standards, information provision, or research collaboration programs, are relevant and necessary (Fagerberg, 2018), we limit our analysis to policies where the budgetary impact can be easily quantified.

In this framework, car manufacturers and users immediately incorporate the policy into their behaviour when making choices. However, we disregard the role of expectations regarding policy stringency (Helm et al., 2003; Aghion, 2019). Policies are treated as exogenous, ignoring the current state of the model, co-evolution of innovation, consumer preferences, and the feasible policy space (Rubin et al., 2015; Wesseling et al., 2015). We assume that the simplified federal and state-level rebate policies from the calibration period continue in place.




Figure 4.3.1: *Simulation of key model outputs from 2001 to 2023. EV uptake and EV sales, in percent (top left); Retail prices for new and used ICEVs and EVs, in dollars (top right); Herfindahl-Hirschman Index, in percent (bottom left); and, average car age, in months (bottom right).*

We consider the interventions listed below:

- **Carbon price:** The gasoline and electricity cost per kilowatt hour increases by $\$ \tau_1$ for each kilogram of CO₂. This policy aims to increase the attractiveness of EVs relative to ICEVs based on their lower emissions per kilometre while increasing the public budget.
- **Subsidy on EV charging electricity:** The electricity cost per kilowatt hour is subsidized by $\tau_2\%$. This policy aims to increase the absolute and relative attractiveness of EVs by decreasing the cost per kilometre. The subsidy only applies to EV charging stations.
- **Rebate for new EVs:** The purchase price of new EVs decreases by $\$ \tau_3$, with a zero-lower bound. This policy aims to increase EV adoption by reducing its sale cost while increasing producers' profits who incorporate the subsidy into their price-setting function.
- **Rebate for used EVs:** The purchase price of used EVs decreases by $\$ \tau_4$, with a zero-lower bound. This policy aims to increase EV sales among car users with high environmental concern and low mobility needs and reduce total CO₂ emissions by minimising the production of new cars.
- **Production subsidy:** The production cost for manufacturing EVs decreases by $\$ \tau_5$, with a zero lower bound. Like adoption subsidy for new cars, this policy is expected to increase EV adoption and producers' profits. However, this policy is expected to induce a lower reduction on prices due to rent seeking.

During the simulated future period, the prices for electricity and gasoline before policy implementation are fixed to the last observed data point. In accordance with the Californian ZEV Plan, we assume a linear decarbonization of the electricity grid, such that by the end of the policy period, emissions per kilowatt-hour of electricity are reduced by 90% by 2035. The rebate on new and used cars implemented during the validation period is removed in order to capture the fact that car sellers enjoy the rebate up to a certain limit of sales.

The policy experiments are evaluated for the period 2024-2035 (144 time steps), according to the ZEV mandate timeline. First, we individually simulate each policy instrument using grid search to vary the policy stringency. The grid search simulates 100 policy values for each instrument, ranging from 0 to the maximum values shown in Table 4.4.1. The maximum values were chosen to represent an extreme value in order to test the limits of each individual policy.

Table 4.4.1: Maximum policy intensity evaluated in the single-instrument scenario.

Instrument	Maximum intensity
Carbon price	1500\$/ TCO_2
Subsidy on electricity	100%
Adoption subsidy for new cars	\$50,000
Adoption subsidy for used cars	\$50,000
Manufacturing cost subsidy	\$50,000

Having evaluated the effect of the policy stringency on multiple measures, we calculate the intensity of the policy required to achieve a target 95% EV uptake on average across simulations, allowing for a 1% discrepancy. We employ a Bayesian optimization model (Head et al., 2021) using a Gaussian process to build a surrogate model that predicts how EV uptake responds to the intensity of the policy. The uncertainty about expected EV uptake decreases with the number of simulation runs, where the selected policy intensity value is chosen by maximizing the expected improvement acquisition function. In doing so, exploration and exploitation are balanced by prioritizing large uncertainty and better fit simultaneously.

The values obtained are used as boundaries to find policy mixes that achieve the target uptake. For each pair of policies, the intensity of one instrument is fixed at 10 equally distributed partitions of the optimal intensity. Using the same Bayesian optimization method, we evaluated the intensity of the policy required by the pairing policy to achieve the target uptake. This analysis excludes pairs that do not include policies that can trigger a transition when implemented in isolation.

Finally, we show the transition path of each policy pair. The intensity chosen for each combination is set to minimize the maximum intensity of the two, relative to its maximal value detailed in Table 4.4.1. This approach is followed by the consideration that large policy instrument intensities can be politically unfeasible or cause limit dynamics in the real world that our model cannot address. Furthermore, some combination of policies can produce positive synergistic effects (Bergh et al., 2021) that we aim to explore. In order to test the stability of the transition, we remove the policy instruments after December 2035 and continue the simulation until December 2050.

4.4.2 Policy results

Individual policies

Figure 4.4.1 shows the effect of each individual policy on EV uptake, net policy cost, cumulative utility, and profits.




Figure 4.4.1: Key model outputs for single policies.

At low values of the carbon price, the EV adoption proportion increases rapidly with a lower marginal increase from approximately \$300/TCO2, achieving a complete transition at around \$1500/TCO2. This high value indicates the magnitude of the behavioral change required for a rapid EV transition. Government revenue increases with the carbon price, but with diminishing returns as a result of higher EV uptake. The manufacturing of new EVs necessary to reach higher uptake responds positively to the stringency of the carbon price, yet with diminishing returns. The production of new EVs increases manufacturers' profits and manufacturing carbon emissions, both cumulative over the period 2025-2035. The increase in manufacturing emissions compensates for the reductions in gasoline use,

leading to higher total cumulative emissions. Furthermore, users' utility, cumulative over the 2025-2035 period, decreases linearly with respect to the value of the carbon price. This is due to the strong negative effects of higher gasoline prices on non-EV-adopters. As a result, despite being the most effective instrument, the carbon price faces severe political feasibility risks reported earlier in the literature (Savin et al., 2024).

The electricity subsidy does not achieve 95% EV uptake for any value. However, EV uptake increases linearly with the electricity subsidy, achieving an uptake of 60% with free electricity. The budgetary impact of the policy is less than other policies due to the high fuel efficiency of EVs and the smaller policy impact on uptake. Given that the policy is insufficient to trigger a substantial production of new EVs, cumulative CO₂ emissions decrease with the size of the subsidy. In this case, the larger manufacturing emissions do not offset the CO₂ savings from lower gasoline use. Furthermore, both cumulative utility and profits linearly increase with the subsidy, the former being more reactive than the latter.

In the case of both the carbon price and the electricity subsidy, the policies affect all car users at all time steps, increasing the impact on utility relative to at-sale policies such as rebates or production subsidies. However, this is counteracted by the high annual discount rate used of 5% (Busse et al., 2013; Greene, 2010), which limits the utility gains EVs have in their long-term behaviour caused by their high efficiency and lower operating costs.

The rebate for new cars increases EV uptake linearly before reaching 95% EV adoption. However, high rebate values causes users to purchase new EVs more often as their price is highly subsidized by the government. As a result, the total cumulative policy cost, the manufacturing emissions, and manufacturers' profits increase more than linearly with respect to the rebate size. In contrast, consumer utility increases linearly and reaches a maximum at the point where the rebate covers the full price of new EVs. These outcomes challenge the feasibility and usefulness of the policy.

In the case of the rebate for used cars, the lack of a pre-existing sizeable stock of used EVs makes the policy highly ineffective, having a negligible impact on any of the outcomes.

The production subsidy has similar effects than the rebate for new cars, since both policies directly affect the price of new EVs. However, its effects are milder as manufacturers are more able to capture rent. Due to higher retail prices, EV uptake does not reach 95% for any production subsidy value. However, the policy-induced uptake at high subsidy values exceeds 90%. The subsidies on manufacturing costs lead to lower turnover of EVs compared to the adoption subsidy for new cars, resulting in lower policy-induced increases in cumulative policy costs, CO₂ emissions, users' utility, and manufacturers' profits.

We calculate the required policy stringency to achieve the uptake target of 95% on average in different stochastic seed simulations. Table 4.4.2 shows the key outputs from the policy intensities compared to the business-as-usual (BAU) scenario.

Our model demonstrates that a 95% EV uptake by 2035 is possible with a carbon price of 910\$/TCO₂ or with a rebate for new EVs of \$36875.57. The suggested value for the carbon price is high relative to some estimates of the social cost of carbon (SCC) but not unheard of (Moore et al., 2024; Bilal and Käenzig, 2024). Likewise, that of the rebate for new cars is substantially larger than the currently implemented rebates worldwide, usually ranging between US\$2500 to US\$20,000 Hardman et al., 2017.

Beyond the difficulty of undergoing a highly stringent intervention, each of the poli-

Table 4.4.2: Outcomes of single policies at minimum intensity achieving a 95% EV fleet

	BAU	Carbon Price	EV Rebate (new)
EV Adoption Proportion, (σ)	0.171 (0.209)	0.954 (0.072)	0.950 (0.046)
Intensity	-	910 \$/ TCO_2	\$36875.57
Cumulative Net Cost, bn USD	0.000	-0.049	0.510
Cumulative Emissions, MTCO₂	120.95	133.53	264.76
Cumulative Emissions (Driving), MTCO₂	99.94	58.87	49.78
Cumulative Emissions (Production), MTCO₂	21.01	74.66	214.98
Cumulative Profit, bn USD	0.025	0.066	0.124
Cumulative Utility (2030), bn USD	3.688	2.410	4.015
Cumulative Utility (2035), bn USD	7.291	6.008	8.197

cies face political feasibility challenges. Mainly, the carbon price induces lower cumulative consumer utility with respect to the BAU scenario, both in the short and in the medium-run, while the rebate for new EVs leads to large policy costs. Furthermore, both policies result in a net increase in cumulative emissions due to the transitional large-scale deployment of EVs. This effect is almost double with the EV rebate. However, the fact that the driving emissions in both cases are lower than in the BAU scenario indicates that the rebound effect is transitional.

As a consequence of the trade-offs highlighted in the outcomes for the single policies studied, we now consider how policy mixes of two instruments may fare better in negating their individual weaknesses. Additionally, we consider whether synergies between policies may help reducing their stringencies to more politically palpable levels.

Policy pairs

This section discusses the implications of combining policy instruments. To this end, we compare the cumulative emissions resulting from each policy combination with their respective cumulative utility and net policy cost in the scatter plots shown in Figure 4.4.2.

The first key observation is that no policy scenario achieves lower cumulative emissions than the BAU scenario by 2035. This is due to the necessary large-scale production of EVs required to reach 95% uptake, which offsets the emissions reduction from lower gasoline use. In the case of pairs including rebates for new EVs, except when combined with a carbon price, these larger emissions are exacerbated due to an overproduction of cars.

The second key observation is that no policy achieves higher utility than in the BAU scenario and a net surplus. However, there are policy combinations that improve consumer utility with respect to the BAU scenario at moderately low policy costs. This is mainly the case for combinations with a small carbon price or for the combination of electricity and production subsidies.

A third observation is that policy combinations achieve substantial reductions in the policy intensity of both instruments, thereby improving the political acceptability of the transition. This is particularly true with policy pairs involving a carbon price, a rebate for new EVs, or a production subsidy.

Figure 4.4.2 also shows how instruments interact with one another in minimising the implied trade-offs exposed in Table 4.4.2. Adding an electricity subsidy to the carbon price can substantially mitigate utility losses, despite not improving on the BAU scenario. Furthermore, by reducing the intensity of the carbon price, this combination also




Figure 4.4.2: *Policy pairs outcomes that achieve 94-6% EV adoption proportion, cumulative from 2024 to 2035. Policy intensity is indicated in half-circle size, where the full circle indicates the maximum intensity. Single instrument policies are shown as a full circle. BAU is shown in black as a reference.*

mitigates the policy-induced increase in manufacturing emissions, achieving the lowest cumulative emissions of all policy pairs. This is because the carbon price initially spurs a first production wave of more efficient ICEVs, and after a second production wave of EVs, once social diffusion dynamics take place with EVs becoming more widely considered. With the introduction of an electricity subsidy, the first wave of efficient ICEVs is minimized due to the lower intensity of the carbon price. Additionally, the positive incentive on EV use driven by the electricity subsidy aids in its early diffusion, thereby avoiding an intermediate wave of new efficient ICEVs. It is also worth noting that these improvements in consumer utility and CO₂ emissions do not cause a net policy cost. The budgetary impact of the electricity subsidy is relatively small due to the high efficiency of EVs.

Combining a carbon price with a rebate for used EVs has similar, yet weaker effects. Supporting the market for used cars leads to reductions in total emissions, relative to solely the carbon price, as it decreases the turnover of new EVs. Furthermore, it leads to higher consumer utility for the users who buy these cars, who are often car users with lower income. With a low intensity of the rebate for used EVs, the cumulative emissions and the net policy cost are minimized due to the lack of rebound effects. However, when compared to the combination with the electricity subsidy, the impact of the rebate for used EVs is worse across the three evaluated dimensions. This is due to both the

constraint stock of used EVs and the one-time utility boost at the purchase decision, whereas the electricity subsidy affects all EV users at all time steps.

The combination of a carbon price and the production subsidy achieves a substantial reduction in the intensity of both instruments. At a relatively low policy cost and rebound effect, adding a production subsidy can fully mitigate or offset the decline in utility caused by the carbon price. This is due to the reduction in the price of EVs and in the carbon price intensity. However, unlike in the case of the pairs commented earlier, adding a production subsidy to the carbon price leads to an increase in emissions due to the higher production of EVs.

A carbon price and new EV rebate leads to substantial decreases in both policy intensities, mitigates or overcomes utility losses, and causes moderately low policy costs and rebound effects. Compared to the combination with the production subsidy, the induced rebound effect and the net policy cost are smaller. This is counter-intuitive because the rebate for new EVs leads to a stronger price reduction than the production subsidy. However, because the effect of the policy on prices is stronger, it allows for a lower intensity of both policies. Therefore, the over-production of efficient ICEVs and new EVs is mitigated.

While the described policy pair examples show how instruments can mitigate the undesired outcomes of the transition by exploiting complementarities, the case of the production subsidy combined with the new EV rebate demonstrates that this is not always the case. When combined, both the cumulative emissions and policy costs substantially increase without improving consumer utility compared to the case where only the new EV rebate is being used. This is because the price for EVs reaches a zero-lower bound, and additional subsidies only boost manufacturers' profits.

The pair of electricity subsidy and rebate for new EVs minimizes the policy-induced rebound effect and the policy cost, while marginally increasing utility. This effect is stronger as the subsidy for electricity increases. The reason for this absolute improvement is that the incentives provided by the electricity subsidy help reducing the intensity of the rebate for new EVs required to meet the adoption targets. As a result, the over-production of new EVs is avoided, resulting in less manufacturing emissions and rebates paid.

As in the case of electricity subsidies, adding a rebate for used cars to the new car rebate decreases costs and emissions while keeping utility equal to solely the new car rebate. More specifically, the rebate for new and used EVs is only effective when the rebate for used cars is small. The rebate for used EVs must be sufficiently large to make some users prefer used EVs over new ones. However, when its intensity is too high, it supports the rebound effect by increasing the overall car turnover.

By allowing for policy pairs, the combination of electricity and production cost subsidies can meet the target uptake despite of not doing so when implemented in isolation. This combination achieves the highest cumulative consumer utility with cumulative emission and policy cost levels akin to those obtained with carbon prices combined with new EV rebates or production subsidies. The performance of this combination is improved in all three dimensions when the intensity of the electricity subsidy is stronger than that of the production subsidy. This is because a lower intensity of the production subsidy mitigates the over-production of EVs, thereby leading to lower manufacturing emissions and subsidies paid. Additionally, subsidizing electricity has a greater impact on utility due to the assumed lack of rent-seeking behaviour by energy providers.

Finally, for each of the 8 policy pairs that achieve the EV adoption target, we identify

the transition path corresponding to the intensity combination that minimizes the maximum relative intensity of the two instruments. Specifically, this combination ensures that the highest relative intensity (as a proportion of its maximum value) between the two instruments is minimized. In Figure 4.4.3 we show the time series of each selected policy for the EV share (top left), the new and used car prices (top right), the flow and cumulative emissions (second row), consumer utility (third row), the average car age (bottom left), and the cumulative net policy cost (bottom right).




Figure 4.4.3: *Policy mix trajectories to 2050 with confidence interval obtained over 64 Monte Carlo simulations.*

We remove the policy instruments in 2035 and run the projection until 2050 in order to verify that the policy-induced transition remains stable. The top left panel shows that regardless of the policy used, EVs remain the preferred option after public intervention stops. Furthermore, the panel shows a slow adoption of EVs under the BAU scenario, where by 2050 the uptake slightly exceeds 40% and the proportion of total sales are 60%. This motivates the need for present-day policy interventions.

The average retail price of EVs converges to a stable value once the policies are

removed. One can observe in the top right panel that the long-run average price of new and used EVs is lower than the BAU scenario for all transition policies. This is caused by two factors. First, the accelerated transition enables more research in EVs, leading to superior technologies. Second, because the entire population uses EVs, firms offer cheaper designs to satisfy low-income customers, whereas in the BAU scenario, EVs are targeted at high-spending segments. Interestingly, in the case where the production subsidy and the rebate for new EVs are implemented, the EV price is higher than in the other policy pairs after the incentives are removed. This indicates that under this policy, innovation seeking cost reductions is not promoted, because the price that car users pay post-rebate is nearly zero, they do not lead to higher profits.

In the long term, all emission flows converge to below the BAU due to the high EV uptake, see second row to the left of Figure 4.4.3. Additionally, we observe that all policies except the combination of production subsidies and the rebate for new EVs achieve lower cumulative emissions than the BAU before 2044. The differences in cumulative emissions are mainly caused by the manufacturing emissions induced during the policy time. After removing the incentives, the flow of emissions is similar in all cases due to the fact that the assumed 90% electricity grid decarbonization reduces the relevance of fuel efficiency in terms of emissions. We observe that the flow of emissions decays at a slower pace when the policy combination includes a production subsidy. This is due to the fact that price adjustment is slower as price adaptation is sluggish. Therefore, the price adjustment after removing the policy takes longer to converge with the other policies, as one can observe in the top right panel of Figure 4.4.3. This effect is proportional to the size of the production subsidy.

During the early policy period, the flow of utility of policy pairs including a carbon price is much lower than that of either the BAU or subsidy policies. Even in the best-case scenario, the yearly utility does not catch up to the BAU until 2031, which may result in a period of sustained political pressure to roll back the carbon price due to its perceived economic burden. Notably, all three combinations of policies that include an electricity subsidy produce a high yearly utility. After 2035, consumer utility gradually declines due to the increasing age of the car fleet. However, after removing the incentives the converging utility flow is larger than in the BAU scenario due to the higher adoption of EVs. Crucially, in the long run the cumulative utility resulting from an EV transition is higher than the BAU scenario in all cases except for the combination of the carbon price and the rebate for used EVs.

In all policy scenarios that achieve the EV uptake target, the car age drops substantially with respect to the BAU scenario due to the renewal of the car fleet. However, after removing the policies, the average car age exceeds that of the BAU in all cases. This lower replacement rate is due to the fact that EVs have much higher efficiencies than their ICEVs counterparts. Thus, car users have little reason to buy new cars when used cars are of a high calibre. In addition to indicating higher satisfaction for EVs, the fact that EVs are kept for longer contributes to sustainability by lowering the manufacturing emissions.

On the one hand the only policy package that achieves a net surplus is the combination of a carbon price with an electricity subsidy. On the other hand, the rebate to new EVs with a production subsidy leads to an undesirable four-fold increase in policy costs. Intermediate costs are found in all other combinations, with pairs including carbon price having lower costs. Notably, in the case of the carbon price and used-car rebate costs are encountered towards the end of the transition as the stockpile of used EVs finally begins

to accumulate, facilitating EV purchases for lower-income car users through the rebate.

From our simulation experiments, we derive several key results. First, policy combinations can improve the political feasibility of the EV transition by reducing the intensities of policies and minimising their induced negative impacts. Counter-intuitively, the addition of supplementary policies to a carbon price can lower the induced emission rebound effect from social locking whilst mitigating medium-run utility losses at a relatively low cost. In the case of the rebate for new cars, some combinations reduce emissions, with respect to the single policy, and policy costs without substantially affecting utility. Furthermore, the pairing of production and electricity subsidies appears as a suitable alternative despite of the fact that neither of the instruments can achieve a 95% EV transition when applied in isolation.

A second takeaway we obtain from our results is that, before 2035, all transition policies entail drawbacks. The required large-scale deployment of EVs causes an increase in manufacturing emissions that exceeds the savings in gasoline use. Furthermore, all policies lead to either losses in cumulative utility or net policy costs. Therefore, policy makers must allocate a budget to the EV transition or face political pushback on implemented policies.

As a third point, in the long run and after removing the incentives, successful transitions remain stable, and the flows of each variable converge. In the aftermath, the transition to EVs leads to a scenario where EVs are cheaper and are kept for longer, consumers are more satisfied, and mobility is less polluting.

4.5 Conclusions

In the pursuit of light transportation decarbonization, this paper proposes an agent-based model (ABM) to examine the complex interplay between consumer adoption, firm innovation, and policy incentives in the transition to electric vehicles (EVs). We build upon previous literature studying EV transitions by analysing policy outcomes in an ABM integrating discrete choice models, social influence dynamics, a used car market, endogenous pricing and directed innovation through an NK framework. By integrating these submodules, we replicate the complex dynamics between car users and manufacturers that are key in the diffusion of innovations. The model is calibrated on the time frame of 2001-2023 in California, with a policy road map considered until 2035. Our study addresses the following research question: what policy combinations, that meet adoption targets can best balance the trade-offs between costs, emissions, and consumer utility?

To answer this question, we simulate a battery of policy instruments to fuel EV adoption, including a carbon price, an electricity and production subsidy, and a rebate for new and used EVs. When implementing single-policy packages, only the carbon price and the rebate for new EVs achieve the 95% target uptake. However, each policy faces severe feasibility challenges due to the high policy intensities required. Moreover, both lead to cumulative emissions greater than the business-as-usual scenario (BAU). In the initial phases of the mobility transition, the scaling of EV production leads to a surge in manufacturing-related emissions, temporarily offsetting reductions from decreased gasoline use. Therefore, we identify the following policy trade-offs: the carbon price generates a surplus but causes a collapse in consumer utility, whilst the rebate for new EVs leads to large policy costs and maximizes the induced emissions rebound.

In order to mitigate these downsides, we explore policy combinations involving two

instruments. By introducing pairs of policies, the intensity and negative outcomes of the carbon price or the new EV rebate can be minimized. Furthermore, the combination of production and electricity subsidies becomes an alternative to meet the uptake target despite their ineffectiveness when implemented in isolation. These results highlight the importance of exploiting instrument complementarity when formulating policy packages in order to reduce rebound effects and short to medium-term trade-offs. Furthermore, we find that not all policy pairs produce positive synergies. In the case of a rebate for new cars and production subsidies, the results are simply greater policy costs and larger emissions with no increase in consumer utility.

In order to test the stability of the EV transition, we simulate the projections until 2050 after removing the incentives in 2035. Once the transition is consolidated, the system converges to a lower flow of emissions and higher utility regardless of the policies implemented to achieve the desired uptake. Ultimately, the widespread diffusion of EVs leads to lower car costs, longer lifespans, increased consumer satisfaction, and reduced pollution from mobility.

Looking ahead, the limitations in the model provide key areas for future research. Firstly, the study of policy combinations in this paper is not exhaustive. Further positive synergies can be found by studying policy mixes including more than two instruments. In addition, other policies may be included in the analysis by improving upon limitations in the model where simplifications have been made.

The expectation formation of users and manufacturers is a fundamental area of future research. Incorporating more sophisticated forecasting techniques and factoring in expected regulation can crucially modify investment and purchase decisions. By doing so, other policies such as standards and regulations or joint research ventures could be assessed within this framework.

Another limitation of our model is the lack of financial and physical constraints. Financial markets are a crucial institution for the large-scale adoption in EVs as they determine the ability of manufacturers to invest in R&D and the required capital for EV production. Physical constraints also play an important role as they limit the speed of EV deployment due to the necessary build-up of productive infrastructure and car stocks. Future research should incorporate capital costs of switching production and a supply model with an explicit representation of production planning and inventory management.

Future research could also incorporate the implications that EV production has on material dependence by endogenizing the cost and emissions associated to the inputs required for EV production. This would allow studying the implications of innovation on material use efficiency as well as the inclusion of a carbon-border adjustment mechanism to price in imported production emissions.

Given that the results of our simulations indicate that emission reductions, relative to the BAU scenario, are only achieved in the long-run, we endorse the study and implementation of alternative policies to minimize car dependence, in both use and ownership. This includes the promotion and improvement of public transport infrastructure, urban re-design aimed at reducing mobility needs, or car sharing schemes to lower car ownership. Moreover, future research should address the price responsiveness car users have in their driving choices. The simultaneous study of urban design to reduce car dependence and price responsiveness in mileage is a promising research avenue as the distance driven becomes more elastic when car mobility needs are minimized.

The absence of sophisticated expectation formation, financial constraints, and supply

chain dynamics means that our analysis does not fully capture real world barriers to rapid EV deployment. Crucially, our findings reinforce that a singular focus on replacing internal combustion engine vehicles with electric vehicles is insufficient to achieve deep decarbonization of transport due to persistent production emissions. Policies must also target the root of the problem: an over-reliance on private car ownership and use.

As a final note, policy makers must apply more stringent policies in order to achieve a quick transition to EVs with cleaner mobility. We strongly encourage exploiting positive synergies between policy instruments to improve the political acceptability of the policy and minimizing rebound effects. However, our study demonstrates that a fast and ambitious reduction in carbon emissions requires tackling the root of the problem: an over-reliance on private car ownership and use.

Code availability statement

The model code and documentation are available at:

www.github.com/danielTorren/endogenous_innovation_and_preference_change

References

Adepetu, Adedamola, Srinivasan Keshav, and Vijay Arya (2016). “An agent-based electric vehicle ecosystem model: San Francisco case study”. In: *Transport Policy* 46, pp. 109–122.

Adner, Ron, Ferenc A. Csaszar, and Peter B. Zemsky (2014). “Positioning on a Multi-Attribute Landscape”. In: *Management Science* 60.11, pp. 2794–2815.

Aghion, P. (2019). “Path dependence, innovation and the economics of climate change”. In: *Handbook on Green Growth*. Ed. by R. Fouquet. Cheltenham, UK: Edward Elgar Publishing, pp. 67–83.

Axsen, Jonn and Kenneth S. Kurani (2013). “Hybrid, plug-in hybrid, or electric—What do car buyers want?” In: *Energy Policy* 61, pp. 532–543.

Baumann, Oliver and Nicolaj Siggelkow (2013). “Dealing with Complexity: Integrated vs. Chunky Search Processes”. In: *Organization Science* 24.1, pp. 116–132.

Bergh, Jeroen van den, J Castro, Stefan Drews, Filippos Exadaktylos, Joël Foramitti, Franziska Klein, Théo Konc, and Ivan Savin (2021). “Designing an effective climate-policy mix: accounting for instrument synergy”. In: *Climate Policy* 21.6, pp. 745–764.

Bilal, Adrien and Diego R Käenzig (2024). *The macroeconomic impact of climate change: Global vs. local temperature*. Tech. rep. National Bureau of Economic Research.

Brown, Maxwell (2013). “Catching the PHEVer: simulating electric vehicle diffusion with an agent-based mixed logit model of vehicle choice”. In: *Journal of Artificial Societies and Social Simulation* 16.2, p. 5.

Busse, Meghan R, Christopher R Knittel, and Florian Zettelmeyer (2013). “Are consumers myopic? Evidence from new and used car purchases”. In: *American Economic Review* 103.1, pp. 220–256.

California Air Resources Board (1990a). *Final Regulation Order. Low-Emission Vehicles and Clean Fuels. California Exhaust Emission Standards and Test Procedures for 1988 and Subsequent Model Passenger Cars, Light-Duty Trucks, and Medium-Duty Vehicles*. Technical Report. California Air Resources Board.

California Air Resources Board (1990b). *Proposed Regulations for Low-Emission Vehicles and Clean Fuels: Final Statement of Reasons*. Technical Report. California Air Resources Board.

California Air Resources Board (1990c). *Transcripts of Public Hearing, September 27 and 28*. Public Hearing Transcript.

Cohen, W. M. and D. A. Levinthal (1990). “Absorptive Capacity: A New Perspective on Learning and Innovation”. In: *Administrative Science Quarterly* 35, pp. 128–152.

Cranmer, Kyle, Johann Brehmer, and Gilles Louppe (2020). “The frontier of simulation-based inference”. In: *Proceedings of the National Academy of Sciences* 117.48, pp. 30055–30062.

Csaszar, Felipe A and Daniel A Levinthal (2016). “Mental representation and the discovery of new strategies”. In: *Strategic Management Journal* 37.12, pp. 2031–2049.

Domarchi, Cristian and Elisabetta Cherchi (2023). “Electric vehicle forecasts: a review of models and methods including diffusion and substitution effects”. In: *Transport reviews* 43.6, pp. 1118–1143.

Dosi, Giovanni, Giorgio Fagiolo, and Andrea Roventini (2010). “Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles”. en. In: *Journal of Economic Dynamics and Control* 34.9, pp. 1748–1767.

Dyer, Joel, Patrick Cannon, J Doyne Farmer, and Sebastian M Schmon (2024). “Black-box Bayesian inference for agent-based models”. In: *Journal of Economic Dynamics and Control* 161, p. 104827.

Eggers, Felix and Fabian Eggers (2011). “Where have all the flowers gone? Forecasting green trends in the automobile industry with a choice-based conjoint adoption model”. In: *Technological Forecasting and Social Change* 78.1, pp. 51–62. ISSN: 0040-1625.

Eppstein, Margaret J, David K Grover, Jeffrey S Marshall, and Donna M Rizzo (2011). “An agent-based model to study market penetration of plug-in hybrid electric vehicles”. In: *Energy Policy* 39.6, pp. 3789–3802.

Fagerberg, Jan (2018). “Mobilizing innovation for sustainability transitions: A comment on transformative innovation policy”. In: *Research Policy* 47, pp. 1568–1576.

Fan, Xia and Chuanju Li (2025). ““To be unfolding” or “be on its last legs”—Preferential tax policies between supply and demand and the development of the new energy vehicle industry based on an ABM model”. In: *Energy Policy* 200, p. 114552.

Feng, Bo, Qiwen Ye, and Brian J. Collins (2019). “A dynamic model of electric vehicle adoption: The role of social commerce in new transportation”. In: *Information & Management* 56.2. Social Commerce and Social Media: Behaviors in the New Service Economy, pp. 196–212. ISSN: 0378-7206.

Ferguson, Mark, Moataz Mohamed, Christopher D. Higgins, Elnaz Abotalebi, and Pavlos Kanaroglou (2018). “How open are Canadian households to electric vehicles? A national latent class choice analysis with willingness-to-pay and metropolitan characterization”. In: *Transportation Research Part D: Transport and Environment* 58, pp. 208–224. ISSN: 1361-9209.

Greene, David L (2010). “How consumers value fuel economy: A literature review”. In.

Greene, David L, Sangsoo Park, and Changzheng Liu (2014). “Analyzing the transition to electric drive vehicles in the US”. In: *Futures* 58, pp. 34–52.

Grieco, Paul L E, Charles Murry, and Ali Yurukoglu (May 2024). “The Evolution of Market Power in the U.S. Automobile Industry”. In: *The Quarterly Journal of Economics* 139.2, pp. 1201–1253.

Hardman, Scott, Amrit Chandan, Gil Tal, and Tom Turrentine (2017). “The effectiveness of financial purchase incentives for battery electric vehicles – A review of the evidence”. In: *Renewable and Sustainable Energy Reviews* 80, pp. 1100–1111. ISSN: 1364-0321.

He, Bao-Jie, Dong Zhao, and Zhonghua Gou (2020). “Integration of Low-Carbon Eco-City, Green Campus and Green Building in China”. In: *Green Building in Developing Countries*. Ed. by Zhonghua Gou. Cham, Switzerland: Springer, pp. 49–78.

Head, Tim, Manoj Kumar, Holger Nahrstaedt, Gilles Louppe, and Iaroslav Shcherbatyi (Oct. 2021). *scikit-optimize/scikit-optimize*. Version v0.9.0.

Helm, D., C. Hepburn, and R. Mash (2003). “Credible carbon policy”. In: *Oxford Review of Economic Policy* 19.3, pp. 438–450.

Herman, Jon and Will Usher (Jan. 2017). “SALib: An open-source Python library for Sensitivity Analysis”. In: *The Journal of Open Source Software* 2.9.

Herrmann, J.K. and Ivan Savin (2017). “Optimal Policy Identification: Insights from the German Electricity Market”. In: *Technological Forecasting & Social Change* 122, pp. 71–90.

Hoen, Anco and Mark J. Koetse (2014). “A choice experiment on alternative fuel vehicle preferences of private car owners in the Netherlands”. In: *Transportation Research Part A: Policy and Practice* 61, pp. 199–215. ISSN: 0965-8564.

Horne, Matt, Mark Jaccard, and Ken Tiedemann (2005). “Improving behavioral realism in hybrid energy-economy models using discrete choice studies of personal transportation decisions”. In: *Energy Economics* 27.1, pp. 59–77.

Hulshof, Daan and Machiel Mulder (2020). “Willingness to Pay for CO2 Emission Reductions in Passenger Car Transport”. In: *Environmental and Resource Economics* 75, pp. 899–929.

Jain, Amit and Bruce Kogut (2014). “Memory and Organizational Evolvability in a Neutral Landscape”. In: *Organization Science* 25.2, pp. 479–493. ISSN: 1047-7039.

Jang, Sungsoon and Jae Young Choi (2021). “Which consumer attributes will act crucial roles for the fast market adoption of electric vehicles?: Estimation on the asymmetrical & heterogeneous consumer preferences on the EVs”. In: *Energy Policy* 156, p. 112469.

Jaramillo, P., S. Kahn Ribeiro, P. Newman, S. Dhar, O. E. Diemuodeke, T. Kajino, D. S. Lee, S. B. Nugroho, X. Ou, A. Hammer Strømman, and J. Whitehead (2022). “Transport”. In: *Climate Change 2022: Mitigation of Climate Change*. Ed. by P. R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, and J. Malley. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK and New York, NY, USA: Cambridge University Press. Chap. 10.

Kauffman, Stuart A (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.

Keyes, Anna K.M. and Douglas Crawford-Brown (2018). “The changing influences on commuting mode choice in urban England under Peak Car: A discrete choice modelling approach”. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 58, pp. 167–176.

Kollmuss, Anja and Julian Agyeman (2002). “Mind the gap: why do people act environmentally and what are the barriers to pro-environmental behavior?” In: *Environmental education research* 8.3, pp. 239–260.

Konc, Théo, Ivan Savin, and Jeroen van den Bergh (2021). “The social multiplier of environmental policy: Application to carbon taxation”. In: *Journal of Environmental Economics and Management* 105, p. 102396.

König, Adrian, Lorenzo Nicoletti, Daniel Schröder, Sebastian Wolff, Adam Waclaw, and Markus Lienkamp (2021). “An overview of parameter and cost for battery electric vehicles”. In: *World Electric Vehicle Journal* 12.1, p. 21.

Kotler, Philip (1989). “From mass marketing to mass customization”. In: *Planning Review* 17.5, pp. 10–47.

Ledna, Catherine, Matteo Muratori, Aaron Brooker, Eric Wood, and David Greene (2022). “How to support EV adoption: Tradeoffs between charging infrastructure investments and vehicle subsidies in California”. In: *Energy Policy* 165, p. 112931.

Lee, Rachel and Solomon Brown (2021). “Evaluating the role of behavior and social class in electric vehicle adoption and charging demands”. In: *Iscience* 24.8.

Levinthal, Daniel A (1997). “Adaptation on rugged landscapes”. In: *Management Science* 43.7, pp. 934–950.

Li, Jingjing, Jianling Jiao, and Yunshu Tang (2019). “An evolutionary analysis on the effect of government policies on electric vehicle diffusion in complex network”. In: *Energy policy* 129, pp. 1–12.

Lynn, Michael (2012). “Segmenting and Targeting Your Market: Strategies and Limitations”. In: *The Cornell School of Hotel Administration on Hospitality: Cutting Edge Thinking and Practice*. John Wiley & Sons, pp. 351–369.

Malerba, Franco, Richard Nelson, Luigi Orsenigo, and Sidney Winter (2008). “Public policies and changing boundaries of firms in a “history-friendly” model of the co-evolution of the computer and semiconductor industries”. In: *Journal of Economic Behavior & Organization* 67.2. Agent-based models for economic policy design, pp. 355–380. ISSN: 0167-2681.

Mazzucato, M. (2013). “Financing innovation: creative destruction vs. destructive creation”. In: *Industrial and Corporate Change* 22.4, pp. 851–867.

Mazzucato, M. and G. Semieniuk (2017). “Public financing of innovation: New questions”. In: *Oxford Review of Economic Policy* 33.1, pp. 24–48.

McCoy, Daire and Seán Lyons (2014). “Consumer preferences and the influence of networks in electric vehicle diffusion: An agent-based microsimulation in Ireland”. In: *Energy Research & Social Science* 3, pp. 89–101.

McFadden, Daniel (1974). “Conditional Logit Analysis of Qualitative Choice Behavior”. In: *Economic Theory and Mathematical Economics*. Ed. by Paul Zarembka. New York, NY: Academic Press, pp. 105–142.

Mehdizadeh, Milad, Trond Nordfjærn, and Christian A Klöckner (2022). “A systematic review of the agent-based modelling/simulation paradigm in mobility transition”. In: *Technological Forecasting and Social Change* 184, p. 122011.

Moore, Frances C, Moritz A Drupp, James Rising, Simon Dietz, Ivan Rudik, and Gernot Wagner (2024). “Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties”. In: *Proceedings of the National Academy of Sciences* 121.52, e2410733121.

Nelson, Richard R and Sidney G Winter (1977). “In search of useful theory of innovation”. In: *Research Policy* 6.1, pp. 36–76.

Ning, Wang, Jiahui Guo, Xiang Liu, and Huizhong Pan (2020). “Incorporating individual preference and network influence on choice behavior of electric vehicle sharing using

agent-based model". In: *International Journal of Sustainable Transportation* 14.12, pp. 917–931.

Noori, Mehdi and Omer Tatari (2016). "Development of an agent-based model for regional market penetration projections of electric vehicles in the United States". In: *Energy* 96, pp. 215–230.

Nuñez-Jimenez, Alejandro, Christof Knoeri, Joern Hoppmann, and Volker H. Hoffmann (2022). "Beyond Innovation and Deployment: Modeling the Impact of Technology-Push and Demand-Pull Policies in Germany's Solar Policy Mix". In: *Research Policy* 51, p. 104585.

Østli, Vegard, Lasse Fridstrøm, Kjell Werner Johansen, and Yin-Yen Tseng (2017). "A generic discrete choice model of automobile purchase". In: *European Transport Research Review* 9.1, p. 16.

Papamakarios, George and Iain Murray (2016). "Fast ε -free inference of simulation models with bayesian conditional density estimation". In: *Advances in neural information processing systems* 29.

Puranam, Phanish, Nils Stieglitz, Magda Osman, and Madan M Pillutla (2015). "Modelling bounded rationality in organizations: Progress and prospects". en. In: *Acad. Manag. Ann.* 9.1, pp. 337–392.

Rivkin, Jan W. (2000). "Imitation of Complex Strategies". In: *Management Science* 46.6, pp. 824–844.

Rogers, Everett M. (2003). *Diffusion of Innovations*. 5th ed. New York: Free Press. ISBN: 978-0743222099.

Rubin, E.S., I.M.L. Azevedo, P. Jaramillo, and S. Yeh (2015). "A review of learning rates for electricity supply technologies". In: *Energy Policy* 86, pp. 198–218.

Savin, Ivan, Stefan Drews, and Jeroen van den Bergh (2024). "Carbon pricing–perceived strengths, weaknesses and knowledge gaps according to a global expert survey". In: *Environmental Research Letters* 19.2, p. 024014.

Schloter, Lukas (2022). "Empirical analysis of the depreciation of electric vehicles compared to gasoline vehicles". In: *Transport Policy* 126, pp. 268–279. ISSN: 0967-070X.

Schmuck, Richard, Ralf Wagner, Günther Höpfl, Tobias Placke, and Martin Winter (2018). "Performance and cost of materials for lithium-based rechargeable automotive batteries". In: *Nature Energy* 3, pp. 267–278.

Shafiei, Ehsan, Hedinn Thorkelsson, Eyjólfur Ingi Ásgeirsson, Brynhildur Davidsdóttir, Marco Raberto, and Hlynur Stefansson (2012). "An agent-based modeling approach to predict the evolution of market share of electric vehicles: a case study from Iceland". In: *Technological Forecasting and Social Change* 79.9, pp. 1638–1653.

Silvia, Chris and Rachel M. Krause (2016). "Assessing the impact of policy interventions on the adoption of plug-in electric vehicles: An agent-based model". In: *Energy Policy* 96, pp. 105–118. ISSN: 0301-4215.

Sobol, Ilya M (2001). "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates". In: *Mathematics and computers in simulation* 55.1-3, pp. 271–280.

Spangher, Lucas, Will Gorman, Gordon Bauer, Yanzhi Xu, and Chris Atkinson (2019). "Quantifying the impact of US electric vehicle sales on light-duty vehicle fleet CO₂ emissions using a novel agent-based simulation". In: *Transportation Research Part D: Transport and Environment* 72, pp. 358–377.

Sun, Xiaohua, Xiaoling Liu, Yun Wang, and Fang Yuan (2019). “The effects of public subsidies on emerging industry: An agent-based model of the electric vehicle industry”. In: *Technological Forecasting and Social Change* 140, pp. 281–295.

Tang, Rui, Simon Fong, Xin-She Yang, and Suash Deb (2012). “Wolf search algorithm with ephemeral memory”. In: *Seventh International Conference on Digital Information Management (ICDIM 2012)*, pp. 165–172.

Tejero-Cantero, Alvaro, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J. Gonçalves, David S. Greenberg, and Jakob H. Macke (2020). “sbi: A toolkit for simulation-based inference”. In: *Journal of Open Source Software* 5.52, p. 2505.

Unruh, G.C. (2000). “Understanding carbon lock-in”. In: *Energy Policy* 28.12, pp. 817–830.

Unruh, G.C. (2002). “Escaping carbon lock-in”. In: *Energy Policy* 30.4, pp. 317–325.

Wang, Yitong, Ruguo Fan, Kang Du, and Xuguang Bao (2023). “Exploring incentives to promote electric vehicles diffusion under subsidy abolition: An evolutionary analysis on multiplex consumer social networks”. In: *Energy* 276, p. 127587.

Watts, Duncan J and Steven H Strogatz (1998). “Collective dynamics of ‘small-world’ networks”. In: *Nature* 393.6684, pp. 440–442.

Wesseling, J.H., J.C.M. Farla, and M.P. Hekkert (2015). “Exploring car manufacturers’ responses to technology-forcing regulation: The case of California’s ZEV mandate”. In: *Environmental Innovation and Societal Transitions* 16, pp. 87–105.

Windrum, Paul, Tommaso Ciarli, and Chris Birchenhall (2009). “Consumer heterogeneity and the development of environmentally friendly technologies”. In: *Technological Forecasting and Social Change* 76.4, pp. 533–551.

Woody, Maxwell, Shawn A Adderly, Rushabh Bohra, and Gregory A Keoleian (2024). “Electric and gasoline vehicle total cost of ownership across US cities”. In: *Journal of Industrial Ecology* 28.2, pp. 194–215.

Zhang, Linlin, Dea van Lierop, and Dick Ettema (2025). “The effect of electric vehicle use on trip frequency and vehicle kilometers traveled (VKT) in the Netherlands”. In: *Transportation Research Part A: Policy and Practice* 192, p. 104325.

Zhang, Qi, Jiangfeng Liu, Kexin Yang, Boyu Liu, and Ge Wang (2022). “Market adoption simulation of electric vehicle based on social network model considering nudge policies”. In: *Energy* 259, p. 124984.

Appendix

4.A Model equations

4.A.1 Car users

The set I collects all car users, each of them being indexed by i .

Car choice

With probability $\eta \in (0, 1)$, car users are activated to consider replacing their car. Given that cars in the second hand market can only be bought once, consumers are activated sequentially in a random order to make car choices and the stock of used cars is updated accordingly.

The alternatives a user considers include the users' own car, all new and used ICEVs on sale, and all new and used EVs if the user considers them. Equation 4.1 defines the set of cars a user considers as alternatives ($A_{i,t}$):

$$A_{i,t} = \begin{cases} \{v \in V_t : v_{v,t} = i \vee v_{v,t} = 0\} \cup Y_t & \text{if } x_{i,t-1} = 1 \\ \{v \in V_t : (v_{v,t} = i \vee v_{v,t} = 0) \wedge z_{v,t} = 0\} \cup \{y \in Y_t : z_{y,t} = 0\} & \text{if } x_{i,t-1} = 0 \end{cases} \quad (4.1)$$

Where the set V_t defines the stock of active cars and the pointer $v_{v,t}$ returns the car users' ID, or 0 if the car is in the market for used cars. The set Y_t contains all new car designs on sale. A crucial difference is that, given that this model assumes production takes place on demand, each new car designs on sale is represented by a vector of attributes and a sale price, but does not represent an actual car. Only when a new car is bought the car of a given design $y \in Y_t$ is produced and included onto the set of active cars V_t with a unique ID v . The binary statements $z_{v,t}$ or $z_{y,t}$ respectively indicate that the car v or the design y corresponds to an EV, while the binary statement $x_{i,t}$ indicates that the user i considers EVs. Note that, with probability $1 - \eta$, the set of considered alternatives only includes the users' own car.

Equation 4.3 describes the utility a user derives from a car alternative:

$$U_{a,i,t} = \beta_i Q_{a,t}^\alpha + \nu (B_{a,t} \omega_{a,t} (1 - \delta_{a,t})^{L_{a,t}})^\zeta - \gamma_i LCE_{a,i,t} - LCC_{a,i,t} \quad (4.2)$$

$$= \begin{cases} -P_{a,t} + \hat{P}_{i,t} - \gamma_i E_{a,t} + \beta_i Q_{a,t}^\alpha + \nu (B_{a,t} \omega_{a,t} (1 - \delta_{a,t})^{L_{a,t}})^\zeta - D_i \frac{(1+r)(1-\delta_{a,t})(c_{a,t} + \gamma_i e_{a,t})}{\omega_{a,t}(r-\delta_{a,t}-r\delta_{a,t})} & \text{if } v_{a,t} \neq i \\ \beta_i Q_{a,t}^\alpha + \nu (B_{a,t} \omega_{a,t} (1 - \delta_{a,t})^{L_{a,t}})^\zeta - D_i \frac{(1+r)(1-\delta_{a,t})(c_{a,t} + \gamma_i e_{a,t})}{\omega_{a,t}(r-\delta_{a,t}-r\delta_{a,t})} & \text{if } v_{a,t} = i \end{cases} \quad (4.3)$$

The parameter β_i represents the idiosyncratic willingness to pay, in dollars, a user assigns to the car quality ($Q_{a,t}$) with diminishing returns scaled by $\alpha > 0$. Similarly, the parameter ν denotes the dollar willingness to pay for car range with diminishing returns scaled by $\zeta > 0$.

The range of a car is expressed as the product of its fuel tank (ICEV) or battery size (EV), denoted by $B_{a,t}$ and expressed in KWh, and its fuel efficiency, denoted by $\omega_{a,t}$ and expressed in kilometres per KWh. We assume a car's fuel efficiency suffers from depreciation at a fixed monthly rate $\delta_{a,t}$, which is specific to the car type (ICEV or EV). In calculating the present fuel efficiency of a car, we express the car age in months with the variable $L_{a,t}$.

The parameter γ_i represents the idiosyncratic willingness to pay, in dollars per kilogram of CO₂, for carbon emissions reductions. In the utility derivation, this multiplies the expected lifecycle emissions of a car, expressed by $LCE_{a,t}$. These can be decomposed by the type-specific manufacturing emissions ($E_{a,t}$), which are only considered if the evaluated car is new, and the expected lifecycle emissions from fuel or gasoline use. These depend on the user's idiosyncratic monthly distance driven, expressed in kilometres per month and denoted by D_i , on the car's fuel efficiency, and on the emissions per kilowatt hour of gasoline (ICEV) or electricity (EV), denoted by $e_{a,t}$, whose expected values into the future equal its most recent value. Emissions resulting from future driving behaviour are discounted in the utility derivation by a monthly discount rate $r > 0$, see Appendix 4.F for derivation.

Similarly, consumers consider the lifecycle costs of a car, which are denoted by $LCC_{a,t}$. The variable $P_{a,t}$ denotes the purchase price of the evaluated car, taking a value of zero if the car is owned by the user. If the user is evaluating a car other than its own, it deducts from the purchase the price at which it can sell its own car, denoted by $\hat{P}_{i,t}$. The user also calculates the discounted cost from gasoline or electricity use, which has the same form as the discounted driving emissions and uses the naive expectations of the cost per kilowatt hour of gasoline or energy, expressed by $c_{a,t}$.

The pointer $\psi_{i,t}$ returns the car chosen by a user from its considered set of alternatives. The probability an alternative is chosen is proportional to its associated derived utility relative to the other alternatives, as described in Equation 4.4:

$$Pr(\psi_{i,t} = a' \in A_{i,t}) = \frac{\exp(\kappa U_{a',i,t})}{\sum_{a=0}^{A_{i,t}} \exp(\kappa U_{a,i,t})} \quad \text{if } v_{a',t} \neq i \quad (4.4)$$

Where $\kappa > 0$ indicates the intensity of choice. A value of $\kappa = \inf$ indicates users are utility maximising, whilst on the other extreme, for $\kappa = 0$, they choose cars randomly with equal probability.

Social imitation

After all car ownership choices have been made, users re-evaluate their willingness to consider an EV. This is governed by the idiosyncratic imitation threshold $\chi_i \in (0, 1)$, expressed as a minimum share of adopter neighbours required for the user to consider adopting. The population subset $H_i \subset I$ collects user i 's neighbours. This set is obtained from a Watts-Strogatz network with the small-world property to replicate word-of-mouth social interactions in highly clustered social circles where individual neighbours are also likely to be connected. We denote with $V_{i,t}^H \subset V_t$ the subset of active cars that belong to car users in i 's network, as defined in Equation 4.5.

$$V_{i,t}^H = \{v \in V_t : v_{v,t} \in H_i\} \quad (4.5)$$

Equation 4.6 uses this set to determine whether an individual considers buying EVs, as a function of the share of EV users in i 's network:

$$x_{i,t} = \begin{cases} 1 & \text{if } \frac{|V_{i,t}^H : z_{v,t}=1|}{|H_i|} \geq \chi_i \\ 0 & \text{otherwise} \end{cases} \quad (4.6)$$

4.A.2 Car supply

The set J collects all firms, each of them being indexed by j .

Market segmentation

Amid consumer heterogeneity in a multi-attribute market, firms segment the market into groups of representative consumers to effectively choose the features of their products and the optimal prices. They do so by building a representative consumer profile for each segment. Each segment $s \in S$ is characterized by preference parameters $(\beta_s, \gamma_s, \nu_s)$ and the binary statement indicating that the segment considers buying EVs (x_s). The preference parameters correspond to the R-tile value of the population, where the parameters R^β, R^γ and R^ν indicate the number of partitions in the parameter space. The population subset $I_{s,t} \subset I$ includes the users whose preferences correspond to those defined in segment s .

Technology choice and pricing

The set of technologies to produce car designs known to a manufacturer is denoted by $M_{j,t}$, which is a subset of all existing technology M defined by the NK landscapes for ICEVs and EVs ($M = M^{ICE} \cup M^{EV}$), see Appendix 4.E.1 for more details. Each technology m is represented by a vector of car attributes and its corresponding coordinates in the NK landscape.

The estimated profit in segment s using a technology m at some price P is denoted by $\Pi_{m,s,t}$ and defined in Equation 4.7:

$$\Pi(P)_{m,s,t} = \frac{(P - C_{m,t})|I_{s,t}| \exp(\kappa U_{m,s,t})}{\exp(\kappa U_{m,s,t}) + W_{s,t}} \quad (4.7)$$

Where $U_{m,s,t}$ is the lifetime utility the car representative consumer with preferences β_s, γ_s , the unit production costs are denoted by $C_{m,t} > 0$ and $W_{s,t}$ indicates the expected state of competition in sector s . This indicator equals the moving average over the last 12 months of the sum of lifetime utilities in the new car market adjusted by the intensity of choice, as described in Equation 4.8:

$$W_{s,t} = \begin{cases} \sum_{t-13}^{t-1} \sum_{m \in Y_t: z_{y,t}=0} \exp(\kappa U_{m,s,t}) / 12 & \text{if } x_s = 0 \\ \sum_{t-13}^{t-1} \sum_{m \in Y_t} \exp(\kappa U_{m,s,t}) / 12 & \text{if } x_s = 1 \end{cases} \quad (4.8)$$

Notice that for segments that serve clients who do not consider EVs, only ICE cars are included in this metric.

For each segment and each known technology, firms define an optimal price ($P_{m,s,t}^*$) that satisfies expected profit maximization, as defined in Equation 4.9, see Appendix 4.G for derivation:

$$P_{s,m,t}^* = C_{m,t} + \frac{1}{\kappa} + \frac{\mathbb{W} \left(\exp(\kappa U_{m,s,t}(P = C_{m,t})) - 1 \right) / W_{s,t}}{\kappa} \quad (4.9)$$

Thus, the optimal profit function may be broken down into 3 components: the production cost, a bounded rationality aspect which tends to 0 as car users become perfectly rational ($\kappa \rightarrow \infty$) and a market advantage component which increases with the superior of a car technology over the state of the past market competition in that segment.

Firms update their supply line and prices with probability $\theta \in (0, 1)$, accounting for frictions in production, price rigidities, or delays in obtaining information on the segment population. When making supply decisions, firms use at most one technology per segment, a simplification justified by strategic considerations to avoid intra-firm competition. Choosing the product mix and the target segments that determine prices is a complex problem that firms resolve

heuristically. Firms rank potential cars within each segment using a car-segment-specific optimal price. They then iteratively select the most profitable car-segment combination, fixing its price across all relevant segments before moving to the next best option. As a result, a car may be produced for multiple segments at a single price if it outperforms other choices. Formally, define a $S \times |M_{j,t}|$ matrix containing the expected segment profit for each known technology ($M_{j,t} \subset M$). Car variants are sequentially added into $Y_{j,t}$ and accordingly priced using the algorithm below:

1. Activate algorithm with probability θ , else end.
2. Denote with m' and s' the column m and the row s corresponding to the cell with the highest expected profit.
3. Include an element y in $Y_{j,t}$ with the features of technology $m \in M_{j,t}$ and the price $P_{s',m',t}^*$.
4. Delete the row s' .
5. For the elements in column m' , lock the price at $P_{s',m',t}^*$ and update the expected profit values accordingly ($\Pi_{s,m',t}(P_{s',m',t})$).
6. Repeat steps 2 to 5 until $|Y_{j,t}| = Y^+$ or until the matrix is empty.

Where $Y^+ > 0$ is a cap on the number of car variants a manufacturer can simultaneously offer. This cap suggests that manufacturers are constraint in their degree of diversification due to strategic considerations such as brand-reputation.

Innovation

With probability $\iota \in (0, 1)$, manufacturers seek innovations to expand the set of known technologies, reflecting that innovation attempts may not be successful in becoming a novel technology (Nelson and Winter, 1977; Dosi et al., 2010). When innovating, manufacturers build a list of candidate technologies to develop ($M_{j,t}^{RD}$). The list of candidate technologies is formed by the unexplored technologies at a local distance (1 Hamming distance from the NK coordinates) with respect to the last explored technology in the ICEV ($m_{j,t}^{ICE} \in M^{ICE}$) and the EV ($m_{j,t}^{EV} \in M^{EV}$) landscapes. Additionaly, the set of candidate technologies includes the current search locations on each landscape to allow firms to not research if the alternative candidate technologies render insufficient profit. Equation 4.10 formally defines the list of candidate technologies:

$$\begin{aligned} M_{j,t}^{RD} = & \{m \in M^{EV} \setminus M_{j,t-1} : \sum_{n=0}^N |\phi_{n,m'} - \phi_{n,m}| = 1\} \\ & \cup \{m \in M^{ICE} \setminus M_{j,t-1} : \sum_{n=0}^N |\phi_{n,m''} - \phi_{n,m}| = 1\} \\ & \cup \{m', m''\} \end{aligned} \tag{4.10}$$

Where $\phi_{n,m}$ defines the coordinates of a technology m on each of the n dimensions of the NK landscape.

The probability of selecting a technology from the list of candidate technologies is proportional to its relative expected profit. The variable $\Pi_{m,t}^*$ returns the expected profit assigned to a technology m optimally priced in its most profitable segment. Denoting the selected researched technology with $m_{j,t}^{RD}$, Equation 4.11 defines its selection probability probability, conditional on probability ι :

$$Pr(m_{j,t}^{RD} = m' \in M_{j,t}^{RD}) = \frac{\exp(\lambda \Pi_{m',j,t}^*)}{\sum_{m=0}^{M_{j,t}^{RD}} \exp(\lambda \Pi_{m,j,t}^*)} \tag{4.11}$$

Where the parameter $\lambda > 0$ indicates the research intensity of choice. The lower the λ , the less likely a firm is to get stuck in a local solution, but also the less likely a firm selects the best solution. This method has been widely used in the literature as a way to balance exploration and exploitation, which is a common issue in complex search problems (Puranam et al., 2015).

Following this step, the search locations in each landscape are updated to the newest car technology researched,

$$\begin{aligned} m_{j,t}^{EV} &= \begin{cases} m_{j,t}^{RD} & \text{if } m_{j,t}^{RD} \in M^{EV} \\ m_{j,t-1}^{EV} & \text{if } m_{j,t}^{RD} \notin M^{EV} \end{cases} \\ m_{j,t}^{ICE} &= \begin{cases} m_{j,t}^{RD} & \text{if } m_{j,t}^{RD} \in M^{ICE} \\ m_{j,t-1}^{ICE} & \text{if } m_{j,t}^{RD} \notin M^{ICE} \end{cases} \end{aligned} \quad (4.12)$$

Finally, the searched technology is added to the set of known technologies. If the firm exceeds its memory capacity ($M^+ > 0$), the oldest technology to be searched not in production is forgotten. Firm's limited memory capacity is a cognitive limitation whereby the searching agent economizes cognitive power by forgetting unused solutions. This limitation has been widely studied in heuristic search algorithms inspired in animal behaviour (Tang et al., 2012). We define a timer $T_{m,j,t}$ in Equation 4.13 for this purpose:

$$T_{m,j,t}(m \in M_{j,t}) = \begin{cases} T_{m,j,t-1} + 1 & \text{if } m \notin Y_{j,t} \\ 0 & \text{if } m = m_{j,t}^{RD} \vee y \in Y_{j,t} \end{cases} \quad (4.13)$$

Hence, the set of known technologies is expressed as in Equation 4.14:

$$M_{j,t} = \begin{cases} M_{j,t-1} & \text{if } m_{j,t}^{RD} \in M_{j,t-1} \\ M_{j,t-1} \cup \{m_{j,t}^{RD}\} \setminus \{\operatorname{argmax}_{m \in M_{j,t-1}} T_{m,j,t}\} & \text{if } |M_{j,t-1}| = M^+ \\ M_{j,t-1} \cup \{m_{j,t}^{RD}\} & \text{if } |M_{j,t-1}| < M^+ \end{cases} \quad (4.14)$$

Used cars

The market for used cars represents the bulk of car purchases in the US², providing an affordable alternative to low-income users.

This model chooses a parsimonious approach for the representation of the market for used cars. Used car dealers are consolidated into a single entity that prices cars based on the market price of the most similar car in the market for new cars, thus implicitly accounting for market segmentation.

When considering buying a car $v \in V_t : v_{v,t} \neq 0$, the second-hand dealer targets a markup $\mu > 0$. Despite the simplification that used car dealers are consolidated into a single entity, the parameter μ serves as a proxy for degree of monopsony in the market. Thus, under perfect competition ($\mu = 0$), the used car dealer buys a car at a price equal to its estimated selling price, denoted by $\mathbb{E}(P_{i,t})$, whereas a perfect monopsony ($\mu = \infty$) buys cars at their scrapping value, which we assume to be a fixed constant $F > 0$. If the estimated sale price of a car is lower than its scrapping value (F), the buying price becomes the scrapping value. Equation 4.15 defines the price the second hand merchant offers for a car v :

$$\hat{P}_{v,t} = \operatorname{MAX} \left(F, \frac{\mathbb{E}(P_{i,t})}{(1 + \mu)} \right) \quad (4.15)$$

For the calculation of the expected market price, the used car dealer uses as a reference the price of the car in the new market with the most similar features. The reference price is denoted by $P_{v,t}^T$ and defined as in Equation 4.17:

$$P_{v,t}^T = P_{y',t} \quad (4.16)$$

²Statista

$$y' = \operatorname{argmin}_{y \in Y_{t-1}} \left| \frac{\omega_{v,t} - \omega_{y,t-1}}{\operatorname{MAX}(\omega_{y,t-1} : y \in Y_{t-1})} \right| + \left| \frac{Q_{v,t} - Q_{y,t-1}}{\operatorname{MAX}(Q_{y,t-1} : y \in Y_{t-1})} \right| + \left| \frac{B_{v,t} - B_{y,t-1}}{\operatorname{MAX}(B_{y,t-1} : y \in Y_{t-1})} \right| \quad (4.17)$$

Where car features are normalized to ensure both dimensions are in the same order of magnitude. The expected market price of a car owned by an individual i is a fraction of its reference price. This fraction is given by its monthly type-specific price depreciation ($\xi_{v,t}$) (Schloter, 2022), as Equation 4.18.

$$\mathbb{E}(P_{v,t}) (v_{v,t} \neq 0) = P_{i,t}^T (1 - \xi_{v,t})^{L_{v,t}} \quad (4.18)$$

The quality and price depreciation monthly rates feature different values because price adjustments do not map one-to-one the quality decay. Given our parsimonious representation of the second-hand market, we parametrize with μ and ξ the off-model dynamics.

Likewise, Equation 4.19 defines the actual sale price of a used car:

$$P_{v,t} (v_{v,t} = 0) = P_{i,t}^T (1 - \xi_{v,t})^{L_{v,t}} \quad (4.19)$$

The used car dealer scraps all cars whose sale price is lower than the scrapping value. Therefore, the set of active cars V_t increases with the purchases of new cars and decreases with the scrapping of old cars, as defined in Equation 4.20³:

$$V_t = V_{t-1} \bigcup_{i=1}^I \{\psi_{i,t} \in Y_t\} \setminus \{v \in V_{t-1} : P_{v,t} < F, v_{v,t} = 0\} \quad (4.20)$$

4.B Model parameters and variables

Table 4.B.1: Parameters

Symbol	Name	Description	Value	Units	Source
η	Car purchase consideration probability	Probability that a car user considers purchasing a car.	8.3	Percent	-
r	Discount rate	Monthly discount rate used to derive the consumer surplus.	0.165	Percent	(Greene, 2010; Busse et al., 2013)
α	Quality diminishing returns factor	Factor controlling the diminishing returns to scale for the contribution of car quality to the total willingness to pay.	0.5	-	-

³To reduce computational cost we impose a maximum capacity of V^+ the second hand market, nonetheless the market size remains an order of magnitude larger than the number of second hand cars sold per time step ensuring ample selection choices, and we scrap all cars that have not been sold after T^+ periods

Table 4.B.1: Parameters

Symbol Name	Description	Value	Units	Source	
ζ	Range diminishing returns factor	Factor controlling the diminishing returns to scale for the contribution of car range to the total willingness to pay.	0.296	-	(Ferguson et al., 2018)
β	Willingness to pay for quality	Monetary value an individual assigns to car quality.	Median 208, 278	USD/Quality	Indirect calibration
ν	Willingness to pay for range	Monetary value an individual assigns to car range.	1173.65	USD/km	(Ferguson et al., 2018)
γ	Willingness to pay for emissions reduction	Monetary value an individual assigns to CO2 emission reductions.	Median 0.066	USD/kgCO2	(Hulshof and Mulder, 2020)
D_i	Distance driven	User-specific monthly distance driven.	Mean 1119.84	km/month	California Energy Commission
χ_i	EV Im- itation threshold	Threshold representing an individual's peer influence in considering EVs.	Beta-distribution (1.1968, 2.6805)	Percent	Indirectly calibrated
κ	Consumer intensity of choice	Parameter indicating the intensity of choice in the logit model.	2.25e-4	-	Indirectly calibrated
ι	Innovation probability	Probability that a manufacturing firm attempts to innovate.	8.3	Percent	-
θ	Production change probability	Probability that a firm updates its supply line and prices.	8.3	Percent	-
λ	Research intensity of choice	Parameter indicating the intensity of choice in the innovation model.	1e-3	-	
M^+	Technology memory capacity	The maximum number of known technologies a firm can remember.	30	Technologies	-
Y^+	Car variant limit	The maximum number of car types a firm can produce.	10	Technologies	-

Table 4.B.1: Parameters

Symbol Name	Description	Value	Units	Source
T^+	Used car time limit	Maximum number of months before a car gets scrapped.	36	Months
V^+	Used car stock limit	Maximum number used cars in the used car market.	300	Cars
μ	Used car dealer markup	Target markup applied by used car dealers when pricing used cars.	100	Percent
F	Car scrapping value	Scraping value of a car.	669.80	USD
ξ	Price depreciation	Monthly type-specific depreciation rate for the price of used cars.	ICEV: 1.16, EV: 0.87	Percent
δ	Fuel efficiency depreciation	Monthly depreciation rate for the car fuel efficiency.	0.175	Percent
N	Number of dimensions in NK landscape	Number of dimensions (components) in the NK landscape representing technologies.	15	Number of components
K	Epistasis	Epistasis in the NK landscape, representing the degree of interdependency between components of the technology.	3	Number of components
E	Manufacturing emissions	Type-specific manufacturing emissions.	ICEV: 10 EV: 14	Tonnes CO2
(C^-, C^+)	Manufacturing costs	Maximum and minimum manufacturing costs in the NK landscape.	(5,7)	Thousand USD
(Q^-, Q^+)	Car quality	Maximum and minimum quality in the NK landscape.	(0,1)	Abstract quality measure
$\rho(C, Q)$	Correlation cost-quality	Correlation coefficient between manufacturing costs and car quality in the NK landscape.	(0)	Correlation coefficient

Table 4.B.1: Parameters

Symbol Name	Description	Value	Units	Source
(ω^-, ω^+)	Fuel efficiency	Maximum and minimum type-specific fuel efficiency in the NK landscape.	ICEV: (0.79,3.09) EV: (2.73,9.73)	km/KWh (Jaramillo et al., 2022, Appendix 10.2)
$\rho(C, \omega)$	Correlation cost-efficiency	Correlation coefficient between manufacturing costs and fuel efficiency in the NK landscape.	(0)	Correlation coefficient
(B^-, B^+)	Fuel tank / Battery size	Maximum and minimum fuel tank (ICEV) or battery size (EV) in the NK landscape.	ICEV: (469.4) EV: (0,150)	KWh (Schmuck et al., 2018)
$\rho(C, B)$	Correlation cost-quality	Correlation coefficient between manufacturing costs and fuel tank or battery size in the NK landscape.	ICEV: 0, EV: .5	Correlation coefficient
C_m	Manufacturing costs	Unit production costs of technology m at time t .	Technology-USD specific	NK model
Q_m	Quality	Catch-all quality measure of car using technology m .	Technology-Abstract specific quality measure	NK model
B_m	Battery/fuel tank size	Size of the battery or fuel tank of car using technology m .	Technology-kWh specific	NK model
ω_m	Fuel efficiency	Fuel or energy efficiency of car using technology m .	Technology-km/kWh specific	NK model
$c_{v,t}$	Fuel cost	Cost per kilowatt hour of gasoline or electricity	Time-variant	USD/KWh ICEV: EIA EV: EIA
$e_{v,t}$	Fuel emissions	CO2 emissions per kilowatt hour of gasoline or electricity	ICEV: 0.331 EV: Time-variant	kgCO2/KWh ICEV: (Jaramillo et al., 2022, Table 10.8) EV: EMBER
J	Number of firms	Number of firms in the simulation.	16	Number of firms (Grieco et al., 2024)
I	Number of car users	Number of car users in the simulation.	3000	Number of car users (Grieco et al., 2024)

Table 4.B.1: Parameters

Symbol Name	Description	Value	Units	Source
R^β	Number of quality-based segments used to create market segments based on the distribution of β_i .	8	-	-
R^γ	Number of environment-based segments used to create market segments based on the distribution of γ_i .	2	-	-
R^ν	Number of range-based segments used to create market segments based on the distribution of ν_i .	1	-	-
ρ_{WS}	Network density	Density connections between car users in Watts-Strogatz network	0.05	-
p_{WS}	Re-wiring probability	Probability of re-wiring close range links to long range links in Watts-Strogatz network	0.1	-

Table 4.B.2: Variables

Symbol Name	Description	Units
Y_t	Set of new cars on sale	Set of new cars on sale by all firms at time t
$Y_{j,t}$	Firm-specific set of new cars on sale	Set of new cars on sale by firm j at time t
V_t	Set of active cars	Set of all active cars owned by users or in the used car market at time t
$A_{i,t}$	Set of considered alternatives	Set of car alternatives considered by individual i at time t .
$z_{v,t}$	Car type	Binary statement specifying if a car v is an electric vehicle.
$v_{v,t}$	Car Owner ID	Identifier of the owner of car v at time t (0 if in the second-hand market).
$U_{a,i,t}$	Consumer utility	Consumer surplus for car alternative a for individual i at time t .
δ	Fuel efficiency depreciation rate	Fuel or energy efficiency monthly depreciation rate.
$L_{a,t}$	Car age	Age of car a at time t .

Table 4.B.2: Variables

Symbol Name	Description	Units
$LCE_{a,i,t}$	Lifecycle emissions of car a for individual i at time t .	kgCO ₂
$LCC_{a,i,t}$	Lifecycle cost of car a for individual i at time t .	USD
$P_{a,t}$	Purchase cost of new or used car a at time t .	USD
$\hat{P}_{i,t}$	Sale price of user i 's car on the second-hand market at time t	USD
$\psi_{i,t}$	Chosen car	ID of individual i 's chosen car at time t .
H_i	Individual i 's network	Set of car users representing i 's network neighbourhood.
$V_{i,t}^H$	Network neighbours cars	Subset of active car that belong to car users in i 's network
$x_{i,t}$	EV consideration	Binary variable indicating whether individual i considers buying EVs at time t .
$M_{j,t}$	Set of known technologies	Set of technologies known to firm j at time t .
$\Pi(P)_{m,s,t}$	Estimated segment-specific profit	Estimated profit a firm expects to make in segment s using technology m at price P at time t .
$W_{s,t}$	Expected segment-specific state of competition	Expected aggregate consumer surplus in sector s at time t .
$P_{s,m,t}^*$	Optimal price	Optimal price of a technology m in segment s at time t
$M_{j,t}^{RD}$	Set of candidate technologies to explore	Set containing all the unexplored neighbouring technologies with respect to the past search locations on each landscape of firm j at time of firm t .
$m_{j,t}^{RD}$	Searched technology	Searched technology by firm j at time t .
$m_{j,t}^{EV}$	Searched location in EV landscape	Searched location in EV landscape by firm j at time t
$m_{j,t}^{ICE}$	Searched location in ICE landscape	Searched location in ICE landscape by firm j at time t
$\vec{\phi}_m$	Technology coordinates	Coordinates of technology m in the NK landscape.
$T_{m,j,t}$	Memory timer	Number of periods at time t since the last time a technology m has been explored or used by firm j .
$P_{v,t}^T$	Reference price	Reference price in the market for new cars to be used as a benchmark when pricing used car v at time t
$I_{s,t}$	Consumer segment set	Set with users belonging to market segment s at time t .

4.C Parameter calibration

This section describes the method chosen to set each of the parameters in the model. Due to the availability of incomplete data, only some of the parameters directly match the real-world data. In order to minimize parameter calibration, we set most of the unobservable parameters at values that are found in previous studies or that lie in a reasonable order of magnitude. For these parameters, we implement the respective sensitivity analyses to fully understand their relevance in the model dynamics. Lastly, the distribution of vehicle users' innovativeness (what proportion of their social network owns an electric vehicle (EV) before they consider purchasing one) is indirectly calibrated using a neural network to match EV uptake in California during 2010 to 2023.

4.C.1 Vehicle choice parameters

Quality scaling exponent (α): The parameter α serves to create a non-linear relationship between quality and willingness to pay, specifically imposing decreasing returns. Due to a lack of data availability, we set this parameter at 0.5 and then perform a subsequent sensitivity analysis.

Willingness to pay for quality (β_i): The parameter β_i indicates the subjective willingness to pay an individual i assigns to a given scaled quality value. Given that quality is measured in abstract units, this value is non-observable. However, after determining the other parameters relevant to vehicle choice, the order of magnitude of this parameter is key to determining the price level. For this reason, we indirectly calibrate this parameter to make the optimal price for a car with average features targeted at the most represented segment equal the mean price reported in the literature (Grieco et al., 2024), in 2020 prices ($P_{m,s} = \$39,290$).

For this estimation, we assume that the average car is an internal combustion engine vehicle (ICEV), and has average production costs, fuel efficiency, and quality. Furthermore, we assume that the cost and emissions per fuel are the last reported values.

From the optimal price derivation, see section 4.G, we know that:

$$\exp(\kappa U(P_{m,s}, \beta_s)) = W_s(\kappa(P_{m,s} - C_m) - 1) \quad (4.21)$$

Taking the log, we get:

$$\kappa U(P_{m,s}, \beta_s) = \ln(W_s(\kappa(P_{m,s} - C_m) - 1)) \quad (4.22)$$

Multiplying by $1/\kappa$:

$$U(P_{m,s}, \beta_s) = \ln(W_s(\kappa(P_{m,s} - C_m) - 1)) (1/\kappa) \quad (4.23)$$

Knowing that from the perspective of the price setter:

$$U(P_{m,s}, \beta_s) = -P_{m,s} - \gamma_s E_m + \beta_s Q_m^\alpha + \nu(B_m \omega_m)^\zeta - \bar{D} \frac{(1+r)(1-\delta_m)(c_m + \gamma_s e_m)}{\omega_m(r - \delta_m - r\delta_m)} \quad (4.24)$$

Re-arranging, β_s is defined as below:

$$\beta_s = \frac{1}{Q_m^\alpha} \left(\ln(W_s(\kappa(P_{m,s} - C_m) - 1)) (1/\kappa) + P_{m,s} + \gamma_s E_m - \nu(B_m \omega_m)^\zeta + \bar{D} \frac{(1+r)(1-\delta_m)(c_m + \gamma_s e_m)}{\omega_m(r - \delta_m - r\delta_m)} \right) \quad (4.25)$$

Thus, we obtain the willingness to pay for quality of the median consumer. Following this step, we use the income distribution data from the [U.S. Bureau of Labor Statistics](#) to assign the values of β_i to the rest of the individuals in such a way that the same proportion is maintained.

Willingness to pay for emission reduction (γ_i): This parameter indicates how much a consumer is willing to pay to reduce one kilogram of CO_2 per kilometre. We know from Hulshof and Mulder (2020) that on average consumers are willing to pay $\$WTP = \46646.65434 (in 2020 dollars) to reduce 1 gram of CO_2 per kilometre. The average consumer is indifferent between two cars where the price of the first is the $\$WTP$ more expensive but saves 1 gram of CO_2 per kilometre. This is shown in the equality below, where we assume $\gamma_i = \bar{\gamma}$ and $D_i = \bar{D}$

$$U_{a,i,t}(\gamma_i = \bar{\gamma}, D_i = \bar{D}, e_{a,t} = e) = +\$WTP + U_{a,i,t}(\gamma_i = \bar{\gamma}, D_i = \bar{D}, e_{a,t} = e + \omega_v) \quad (4.26)$$

$$(4.27)$$

Then, we can isolate the willingness to pay, and γ_i as follows:

$$\$WTP = \frac{\bar{\gamma}\bar{D}(1+r)(1-\delta_a)}{(r-\delta_a-r\delta_a)} \quad (4.28)$$

$$\bar{\gamma} = \frac{\$WTP(r-\delta_a-r\delta_a)}{\bar{D}(1+r)(1-\delta_a)} \quad (4.29)$$

The willingness to pay for emission reduction γ is normally distributed across the population with standard deviation $\sigma = 39160.31118$ (Hulshof and Mulder, 2020). We assume the distribution of this value is uncorrelated with the other preference parameters, such as willingness to pay for quality or range.

Range scaling exponent (ζ): The parameter ζ serves to create a non-linear relationship between range and willingness to pay, specifically imposing decreasing returns. Ferguson et al. (2018) report that people's willingness to pay per kilometre range is lower for gasoline cars. Potentially, this fact is attributed to the diminishing returns to scale with respect to range (Hoen and Koetse, 2014). Having the range for both cars (700 km ICE, 250 km EV) and the respective WTP/km (\$10.09 ICE, \$20.83 EV)⁴, we proceed to calculate the value of ζ that solves the system of equations below:

$$\$10.09 \times 700 = \nu 700^\zeta \quad (4.30)$$

$$\$20.83 \times 250 = \nu 250^\zeta \quad (4.31)$$

$$(4.32)$$

Which gives the following solution:

$$\zeta = \frac{\ln (\$10.09 \times 700) - \ln (\$20.83 \times 250)}{\ln (700) - \ln (250)} \quad (4.33)$$

$$\zeta = 0.296 \quad (4.34)$$

Willingness to pay for range capacity (ν): This parameter indicates how much a consumer is willing to pay for an extra kilometre in car range. Having calibrated ζ to match the points in previous empirical studies (Ferguson et al., 2018), we solve for ν in the system below:

$$\$10.09 \times 700 = \nu 700^{0.2960778787} \quad (4.35)$$

$$\$20.83 \times 250 = \nu 250^{0.2960778787} \quad (4.36)$$

$$(4.37)$$

Isolating ν :

$$\nu = \$10.09 \times 700 / 700^{0.2960778787} = \$1,015.54 \quad (4.38)$$

$$\nu = \$20.83 \times 250 / 250^{0.2960778787} = \$1,015.54 \quad (4.39)$$

$$(4.40)$$

⁴The survey reports heterogeneity by customer type. We take the weighted average of each WTP to reflect a single representative consumer type

Since this study uses 2015 data, we translate it to 2020 prices using the California CPI [State of California Department of Finance](#), giving us a value of \$1,173.65.

Distance travelled (D_i): This parameter scales a consumer's willingness to pay for savings in fuel costs and fuel emissions. We assign the parameter values to the different drivers in our model based on surveys using self-reported data about their yearly mileage from the [California Energy Commission](#). We simplify in the following fashion: First, we assume that distance is equally distributed over the year. Second, we only take the distribution of distances travelled for new cars, even though they slightly decrease as cars get older. Third, given that the survey data is distributed in ranges, we assign the mid-point to each consumer belonging to a bin. Fourth, given that the last range includes people who drive more than 25,000 miles per year, we assume that the range midpoint is 1.5 times that value. Finally, we assume the distribution of monthly mileage is uncorrelated to other preference characteristics.

Discount rate (r): This parameter influences the relevance of making savings in the future from considering cars with greater efficiency, or EVs when the price and emissions of electricity are lower than those of gasoline. We use a value of 5% yearly discount rate, which gives a monthly value of 0.16515813%. This rate is common in other empirical and modelling literature related to car purchases (Woody et al., 2024; Busse et al., 2013; Greene, 2010), with is value corresponding to near average values of car loan APRs.

Consumer intensity of choice (κ): This parameter determines the degree of rationality in consumer choice, ultimately affecting the speed of EV uptake when EVs become superior to ICE cars. Furthermore, it affects prices as firms optimally levy higher mark-ups when consumers feature bounded rationality. Given that the logit choice model weights the absolute consumer surplus, the value of κ alone is not informative about the degree of rationality consumers feature. It is set through a grid search that balances prices, market concentration and firm markups to match those found in Grieco et al. (2024).

4.C.2 Social network parameters

Network density and re-wiring: To represent the medium over which vehicle users interact a Stogatz-Watts small-world network (Watts and Strogatz, 1998) is employed. In this network, neighbours are highly clustered triads whereby two neighbours of an individual are likely to be neighbours. The network is generated by first forming a ring of vehicle users each connected to their nearest neighbours, such that node degree is homogenous, then with a given probability p_{WS} one end of link between vehicle users is moved to another random user. This re-wiring produced the small world property whereby the network has a short path length due to a few cross-network connections. To test different network sizes, a network density is set ρ_{WS} , instead of setting a fixed number of connections per vehicle user. Thus the mean number of connections per vehicle user is given by $N_{SW} = (I - 1)\rho_{WS}$. The network density is a key in determining how sparse connections are, making it more likely to have clusters of non-adopters whenever density is low. Due to a lack of data availability, we fix this parameter at $\rho_{WS} = 0.05$ and perform a subsequent sensitivity analysis.

Degree of innovativeness ($\chi : \alpha_\chi, \beta_\chi$): This parameter determines the threshold of EV adopters in someone's network before the individual considers buying an EV. It is crucial to determine the speed of EV uptake once EVs begin to outperform ICEVs. We represent the distribution of χ values using a beta distribution. Due to its relevance and lack of data availability, the two parameters that characterise the innovativeness distribution α_χ, β_χ using simulation-based inference (Papamakarios and Murray, 2016; Cranmer et al., 2020; Dyer et al., 2024) to match EV uptake during the validation period. Specifically, we use the neural poster estimation method from the *sbi* python package (Tejero-Cantero et al., 2020). In this method, the outputs of the ABM are used to train a mixture density network that estimates a posterior for the values of α_χ, β_χ . Each simulation run for a given parameter setup, generates the output

vector of yearly EV uptake, in October, during the validation period, these are compared to the real-world uptake data. After an initial round of simulation the posterior estimation generated by the mixture density network can generate new prior values of α_χ, β_χ , narrowing down the range to be explored in the next round. Thereby generating more accurate predictions in a computationally efficient manner. Once training rounds are completed, the posterior is sampled to indicate what values α_χ, β_χ give the best fit, 4.H.2. The best suggested values are tested, with a priority given to those that better match more recent EV adoption data.

4.C.3 Exogenous car parameters

Gasoline and electricity cost per kilowatt hour ($c_{v,t}$): These exogenous variables are type-specific and influence the relative attractiveness of EVs, both in relative terms and in absolute terms, given that EVs have higher fuel efficiency. For ICE cars, we take the monthly cost per gallon in California reported by the EIA [EIA](#) and use the kWh equivalent per gasoline gallon [AFDC](#). For EVs, we use the monthly price of residential electricity per kWh in California [EIA](#). For months with incomplete data, we assume the value equals that of the closest month with reported values. Figure 4.H.3 shows the evolution of both variables.

Gasoline and electricity emissions per kilowatt hour ($e_{v,t}$): These exogenous variables are type-specific and influence the relative attractiveness of EVs, both in relative terms and in absolute terms, given that EVs have higher fuel efficiency. For ICE cars, we take the assumed value in the 2022 IPCC Mitigation report (Jaramillo et al., 2022, Table 10.8) and transform it to kg of CO_2 per kWh. For EVs, we use the monthly data for average emissions intensity in the Californian electricity grid [EMBER](#). For months with incomplete data, we assume the value equals that of the closest month with reported values. Figure 4.H.4 shows the evolution of both variables.

ICE car and EV manufacturing emissions ($E_{v,t}$): These parameters influence the willingness to pay for new cars of environmentally-concerned consumers. We use the production emissions estimated by the IPCC 2022 mitigation report (Jaramillo et al., 2022, p. 1076), where the manufacturing emissions of EVs are estimated at 14 tonnes of CO_2 and those of ICE cars are estimated at 10 tonnes of CO_2 .

ICE car and EV fuel efficiency depreciation rate (δ): These parameters influence the speed at which consumers change their car and diminish the long-term advantages of cars with higher fuel efficiency. Under the constraint that $\delta < \frac{r}{1+r}$, we know the value is upper bounded at 0.246%. Given that, we fix the monthly depreciation rate for both car types at 0.1% and perform a subsequent sensitivity analysis. This value is set so that the average age of vehicles matches that of our stylized facts (10-12 years old), see [Bureau of Transportation Statistics](#).

Scrap value (F): The scrap value increases the speed at which consumers change their car as it increases the minimum price paid by the second-hand merchant, and therefore decreases the net cost from buying a new or second-hand car. We assume a scrap value of \$669.80, which is the reported value in a specialised source [JunkCarsUS](#).

Price depreciation rate for second-hand cars ($\xi_{v,t}$): The price depreciation rate for second-hand EVs and ICE cars has an ambiguous effect. On the one hand, it increases the probability that a second-hand car is bought as it reduces its selling price. On the other hand, it reduces the number of periods before it gets disposed. Its effect on the overall sale of used cars depends on its interaction with the fuel efficiency depreciation rate and the scrap value. We set these parameters based on previous studies (Schloter, 2022), we set a monthly depreciation rate of .87% for ICE cars and one of 1.16% for EVs.

Used car dealer mark-up (μ): This parameter affects the speed at which price-sensitive vehicle users change their car, given that larger values of μ decrease the revenues from selling a user's own car. Lacking reliable information on the actual target mark-up for used cars, which is heterogeneous across used car dealers, we fix this parameter at 100% and perform the

subsequent sensitivity analysis afterwards.

4.C.4 NK landscape parameters

Number of components (N): The number of components affect the performance distribution across technologies. Given that the performance of an attribute equals the average of N draws from a uniform distribution, the distribution of performances over the 2^N existing technologies approaches the normal distribution as N grows larger. For this reason, we fix $N = 15$ as a number sufficiently large to ensure that property. We assume the NK landscapes for ICE cars and EVs are equally large.

Landscape complexity (K): The degree of complexity increases the number of local solutions in the search process, hence decreasing the probability a local-searching agent finds a global solution. As K approaches $N - 1$, the search process becomes closer to a random draw from a normal distribution. Since the purpose of incorporating NK landscapes into this model is to integrate the role of structured search in determining path-dependent industrial transformation in the context of multiple attributes, we select an intermediate level of complexity that ensures the existence of "local traps" while maintaining past search relevance. We fix $K = 3$ and perform the subsequent sensitivity analysis afterwards. We assume the NK landscapes for ICE cars and EVs are equally complex.

Attribute performance correlations (ρ): This vector of parameters determines how correlated are the performance of the difference attributes, potentially creating trade-offs when choosing research paths. To account for the main cost driver in EV manufacturing, we set a positive correlation between manufacturing costs and battery size of 0.5. We assume that the correlation of production costs with quality and fuel efficiency are zero. In doing so, despite of setting the boundaries of each attribute as the currently existing best technology, the model allows for radical innovations by enabling technologies that render the existence of high-quality and efficiency cars at a low cost.

Quality upper and lower bounds in the NK landscape (Q^-, Q^+): Quality is a catch-all variable for all other relevant features that determine the price of a car, such as comfort, style, horsepower, or other appliances (Østli et al., 2017; Grieco et al., 2024). Since this combined measure cannot be observed in real-world data, we create an abstract score that scales in the range (0, 1). Furthermore, we assume that these boundaries are identical in the ICE and EV landscapes, as the main differences, such as battery autonomy, fuel tank and fuel efficiency, are already included in our analysis.

Production costs upper and lower bounds in the NK landscape (C^-, C^+): In setting production boundaries, we represent possible values across a very broad timeline: during the burn-in period, calibration period, policy period and long-term trends. Therefore, we cannot find adequate data sources as our NK landscape does not allow for radical innovation that changes these boundaries. Instead, we take reference order of magnitude values for production costs, sales price, and car markups from literature (König et al., 2021; Woody et al., 2024; Grieco et al., 2024) and then perform grid search until the vehicle prices match those found in our stylized facts during our calibration period.

Fuel tank and battery capacity bounds in the NK landscape (B^-, B^+): We assume that fuel tank size has reached technological maturity and all cars have the same value. Therefore, we fixed it at 469.4 kWh, after converting the fuel tank size of the four most sold cars in California in 2020, see [California Auto Outlook](#), and their associated fuel efficiency, then converting gallons of gasoline to kWh. The battery capacity for EVs is heterogeneous and can benefit from innovation. Based on (Schmuck et al., 2018), we set the maximum value for battery capacity and apply the adjustment described in section 4.E.2 to obtain the values for the NK landscape boundaries. Fixing the lower bound at 0, we reach an upper bound of 150 kWh.

Fuel efficiency bounds in the NK landscape (ω^-, ω^+): We take the values from the

IPCC 2022 Mitigation report (Jaramillo et al., 2022, Appendix 10.2) and use the adjustments described in section 4.E.2 to obtain the values for the NK landscape boundaries. The observed fuel efficiency values for gasoline vehicles range from 1.25 to 2.62 km/KWh⁵, leading to lower and upper bounds in the NK landscape of .791 and 3.087. For EVs, the IPCC reports values efficiency ranges of 4.132 to 8.333, leading to lower and upper bounds in the NK landscape of 2.731 and 9.734.

4.C.5 Car manufacturer parameters

Segment partitions (R): The number of segment partitions on each segment influences the degree of heterogeneity in products in the market and emphasises the path-dependent nature of firm specialisation. Together with the parameters that influence the path-dependence of search (K, λ), this vector of parameters ultimately affects market concentration. We fix the number of partitions at 8 in the β (quality/luxury) space, 2 in the γ (environment) space, and 1 in the ν (range) space. For robustness, we perform the subsequent sensitivity analyses. The high segmentation in quality space allows for more diversity in car models, reflecting low concentration in real-world markets.

Research intensity of choice (λ): This parameter balances exploitation and exploration during the search process. For high values of landscape complexity (K), a low value of (λ), relative to the order of magnitude of expected profits, helps firms escaping local solutions. We set this value at $1e - 5$ and perform the subsequent sensitivity analyses.

Firm memory size (M^+): As in the research intensity of choice (λ), there is a trade-off in setting the firm memory size. A large memory size contributes to research exploitation by avoiding the repetition of research efforts. However, if an agent has all neighbouring technologies in memory it remains permanently stuck until one of them is forgotten. We set the memory size at the exact amount to make possible, yet highly unlikely, a permanent stuck position $M^+ = 30$, given that each of the two landscapes has 15 neighbouring points in every position.

4.C.6 Event-based probabilities

The probability that a vehicle users considers changing its car (η), a car manufacturer consider changing its supply line and prices (θ), and the probability a car manufacturer succeeds at innovating (ι) are set in such way that, on average, these events occur once per year (1/12).

4.D Sensitivity analysis

Figure 4.H.6 shows the sensitivity of key model parameters in the simulated EV uptake. We test the sensitivity of the quality scaling component (panel A), the discount rate (panel B), the used car dealer mark-up (panel C), the consumer intensity of choice (panel D), the parameters governing the chi distribution of the imitation threshold (panels E and F), the research intensity of choice (panel G), the fuel efficiency depreciation rate (panel H), and the technology complexity in the EV and ICEV landscapes (panels I and J).

The quality scaling component negatively impacts EV adoption as it increases the relevance of attributes that are superior among existing ICEVs due to research advantage developed during the burn-in period. The discount rate negatively impacts EV adoption as it diminishes the utility contribution of the long-run benefits of adopting an EV. These benefits are the monetary and emissions savings derived from better fuel efficiency. The mark-up in the market

⁵We converted the units of the IPCC report using a factor of 0.278 KWh per MJ

for used cars has a small positive effect on EV adoption as it makes the long-term aspects of car evaluation relatively more important than its short-term counterparts.

The consumer intensity of choice has a positive effect on EV adoption as consumers become more likely to adopt EVs when they become a superior alternative to ICE cars. EV uptake increases when the distribution of the imitation threshold is left-skewed, favouring social diffusion. A left-skewed distribution corresponds to high values of b_χ and low values of a_χ .

In the case of the research intensity of choice, we find that EV uptake increases when research choices become more random. This is caused by the fact that the policy incentives favouring EV penetration become more relevant than the initial advantage that ICEVs gain over EVs during the burn-in period via research. Additionally, when the pool of potential EV users is small a more stochastic approach to research leads to more EV designs chosen, despite of lower expected profit. The fuel efficiency depreciation rate negatively affects EV uptake because fuel efficiency is the main advantage that EVs have over ICEVs.

The EV landscape complexity has a negligible effect on EV uptake, due to a small numbers of EV sold during the calibration period. Conversely, the ICE landscape complexity has a mixed but significant effect on the EV uptake. On the one hand, with high complexity the likelihood of finding good innovations decreases due to a higher number of local solutions. This makes innovation in the EV landscape more effective. On the other hand, EV uptake increases when the complexity in the ICEV landscape is low as it reduces the variety of technologies that are developed. Lastly, when stochastic seeds are varied to give a confidence interval we assume that the inputs to the model are fixed, thus we only test changes in the dynamics of the model. However, by varying the complexity of the landscape, we change the inputs. This means that differences between complexity runs may also be due to different seeds in the landscape generation.

To complement these local sensitivity analyzes, we additionally perform a Sobol sensitivity analysis (Sobol, 2001), using the SALib Python library (Herman and Usher, 2017), see Figure 4.H.7 and Figure 4.H.8.

The simplest outputs to explain are those of EV adoption proportion, market concentration, utility and mean car age. These have one or two key parameters and high first-order Sobol values, indicating limited interaction between parameters.

In the case of EV adoption, it is dominated by the a_χ that forms one-half of defining the beta distribution of innovativeness thresholds in the model. Specifically, a larger value of a_χ results in a greater threshold and, thus, more social lock-in to prevent EV adoption. We also note the small role that the discount rate r has in EV adoption due to greater values of this parameter, which makes the long-term benefits of EVs less attractive and thus lowers EV adoption likelihood.

The beta distribution parameter a_χ also dominates the market concentration. HHI remains relatively stable with no EV transition due to the long burn-in period; however, if an EV transition occurs, HHI increases rapidly, resulting in greater variance over different a_χ that determines EV adoption rates.

In the case of utility, the quality diminishing returns factor, α , generates more utility for the same car if it has a higher value. However, of greater influence is the consumer intensity of choice, κ , the greater its value, the more individuals maximise utility, but additionally, the lower firm profits mean a corresponding drop in vehicle prices.

In the case of the mean car age, the key parameter is the car efficiency decay δ ; the greater its value, the faster the vehicle fleet has to be replaced, thereby lowering the car age. Additionally, if an EV transition occurs, then vehicle users purchase cars ahead of schedule, due to the efficiency and lower running costs of EVs, thereby lowering the car age of the fleet, hence the importance of a_χ .

In the case of cumulative emissions and firm profit, we see a divergence between the importance of parameters between the Sobol first and total order indices. Both outputs are correlated

in part because the sale of a vehicle results in large production emissions and firm profits.

In the case of cumulative emissions parameters, determining EV adoption plays a key role, but additionally, those such as the used car deal markup are also of importance. In the case of this markup, the greater the value, the lower the trade-in value of used cars is, thus the lower the incentives for individuals to change their cars, resulting in less frequent purchasing of vehicles (either new or used). This change in purchase frequency may have an indirect effect on the production emissions as those purchasing used cars may do so more often, whilst the purchase frequency of new cars remains fixed (due to lower relative price sensitivity). This would also account for the lack of influence of μ on the mean car age, and the effect of μ . In the case of the Epistasis parameter for ICEVs, K_{ICE} , greater landscape complexity may result in lower production emissions as consumers do not see significant technological advances between vehicle purchasing opportunities.

In the case of firm profits, the consumer intensity of choice κ is a key parameter due to greater values driving down prices. Second, greater values of the fuel efficiency decay rate δ result in more frequent purchasing behaviour of new cars as the utility of used hand cars decay more quickly.

4.E The technology landscape

4.E.1 How NK landscapes work

The development of car types is the result of combining different inputs such as frame components, engines, transmission components, batteries, interior and exterior components. From each component, the car producer has a variety of types, shapes or materials to choose from. The choice of which input type to use for the manufacturing of a car has direct and indirect effects on the overall attributes of the final good, provided that the components of a vehicle often complement each other.

We employ NK models (Kauffman, 1993; Levinthal, 1997) to capture the complexity inherent in product design choices. An NK landscape maps all possible input combinations to a fitness outcome. In a stylized and deterministic manner, the modeller can adjust the extent to which the fitness of a specific input combination can be inferred from a sufficiently large record of alternative fitness combinations. This is achieved by modifying the interdependence among inputs, represented by the parameter K . This parameter defines the number of inputs whose states influence the performance contribution of a given input. In other words, K serves as a proxy for the underlying complexity of the problem represented in the NK model.

In our model, the fitness values in the NK landscape represent features of the vehicle that ultimately determine the manufacturer's profits for a given input combination. These features include production costs, vehicle quality, and fuel efficiency. As in prior applications of NK models (Adner et al., 2014; Csaszar and Levinthal, 2016), we represent fitness as a vector of attributes. This approach is justified by evidence (Horne et al., 2005; Østli et al., 2017) that consumers consider multiple factors when selecting a vehicle.

An NK landscape is a discrete hypercube where each vertex represents a combination of inputs in specific states. Let N be the set of components. Each component, indexed by $n \in 0, \dots, N-1$, assumes a discrete state ϕ_n from $\Phi \in \mathbb{Z}$ possible states. As is standard in NK models, we reduce the number of states to a binary choice. Each unique combination of input states ($\vec{\phi}_{n,m}$) corresponds to a specific variant $m \in M$. The index m is derived by converting the binary string of input states to a decimal representation.

The NK model assigns a vector of feature values to each combination of input states. In its generalized form, these feature values are denoted as $\vec{o}_{f,m}$. Each component of this vector corresponds to a normalized fitness value $o_{f,m} \in (0, 1)$ for a specific attribute f . Following

(Adner et al., 2014), we generate correlated distributions of feature values, with the parameter $\rho_f \in [-1, 1]$ controlling the correlation of each feature f relative to a baseline feature $f = 0$. The fitness distribution for $f = 0$ is defined first with $\rho_0 = 1$ by construction.

For each component n , we define a subset $N_n \subset N$ that includes component n and K additional components. This subset determines the inputs whose state combinations influence the performance contribution of n across all features. Interdependent components are typically selected sequentially, but alternative structures, such as randomized or clustered interdependence, can be adopted depending on the application.

We construct a "performance matrix" with $|N|$ columns and 2^{K+1} rows. Each cell in this matrix is assigned an independent random draw from a distribution $\mathbb{D} \in (0, 1)$. This value represents the normalized performance contribution of a component n to the baseline feature $f = 0$, conditional on the specific input state combination in N_n .

To generate fitness contributions for other features ($f \neq 0$), we assign each feature a subset $N_f \subset N$ of $\lfloor \rho_f |N| \rfloor$ components. For components in this subset, the fitness contribution for the characteristic f matches that for $f = 0$ if $\rho_f > 0$ and is equal to $1 - o_{0,n,m}$ if $\rho_f < 0$. For components not in N_f , the contribution is independently drawn from \mathbb{D} . The performance matrix thus contains a vector of fitness contributions for each input combination across all features.

To compute the total fitness vector for a technology $m \in M$, we average the performance contributions of its components for each feature. This is achieved using the row in the performance matrix corresponding to the binary to decimal transformation of the input state combinations in N_n . Finally, the absolute performance values for each feature $f \in F$ are normalized by Min-Max scaling, as shown in Equation 4.41:

$$o_{f,m} = o_f^- + \left(o_f^+ - o_f^- \right) \frac{1}{|N|} \sum_{n=0}^{N_z} o_{n,f,m} \quad (4.41)$$

Where Y_f^- and Y_f^+ are the lower and upper bounds of feature f .

4.E.2 NK landscape properties

The NK landscape generates a range of values rather than a specific value in itself. We can set the upper and lower bounds of each attribute's performance contributions. However, given that the actual performance is the average of N iid draws, the actual performance is not likely to be the input parameter. Figure 4.H.5 shows some properties of NK landscapes for $N = 15$ (a large N is required to generate normally distributed values). We find that the normalized performance values are in the range (.2, .8) with mean and median of 0.5.

We want to find the boundaries (F^-, F^+) in the NK landscape that give us Q^-, Q^+ , knowing that:

$$\begin{aligned} Q^- &= F^- + 0.2(F^+ - F^-) \\ Q^+ &= F^- + 0.8(F^+ - F^-) \end{aligned} \quad (4.42)$$

Subtracting both sides we get:

$$Q^+ - Q^- = 0.6(F^+ - F^-) \quad (4.43)$$

Then we can isolate:

$$F^+ - F^- = (Q^+ - Q^-)/0.6 \quad (4.44)$$

Substituting in both equations we get:

$$\begin{aligned} Q^- &= F^- + 1/3(Q^+ - Q^-) \\ Q^+ &= F^- + 4/3(Q^+ - Q^-) \end{aligned} \quad (4.45)$$

Isolating F^- we get:

$$F^- = (4/3)Q^- - (1/3)Q^+ \quad (4.46)$$

Substituting F^- we can isolate F^+ as:

$$F^+ = Q^+(4/3) - Q^-(1/3) \quad (4.47)$$

As a general expression, where π^+, π^- are the respective percentages we are targeting:

$$\begin{aligned} F^- &= \frac{Q^-\pi^+ - Q^+\pi^-}{\pi^+ - \pi^-} \\ F^+ &= \frac{Q^+(1 - \pi^-) - Q^-(1 - \pi^+)}{\pi^+ - \pi^-} \end{aligned} \quad (4.48)$$

4.F Lifecycle evaluation derivation

For an infinite geometric series $\sum_{t=0}^{\infty} Yx^t$, where $|x| < 1$, the sum is $\frac{Y}{1-x}$. With no loss of generality, we ignore the case where the user is evaluating its own car. We make the following notation simplifications:

$$C = E_{a,t} \vee P_{a,t} - \hat{P}_{i,t} \quad (4.49)$$

$$z = c_{a,t} \vee e_{a,t} \quad (4.50)$$

$$LCC_{a,i,t} = C + \sum_t^{\infty} D_i \frac{z}{\omega_{a,t}(1+r)^t(1-\delta_{a,t})^t} \quad (4.51)$$

$$= C + \frac{D_i z}{\omega_{a,t}} \sum_t^{\infty} \left(\frac{1}{(1+r)(1-\delta_{a,t})} \right)^t \quad (4.52)$$

$$= C + \frac{\frac{D_i z}{\omega_{a,t}}}{1 - \frac{1}{(1+r)(1-\delta_{a,t})}} \quad (4.53)$$

$$= C + \frac{\frac{D_i z}{\omega_{a,t}}}{\frac{(1+r)(1-\delta_{a,t})-1}{(1+r)(1-\delta_{a,t})}} \quad (4.54)$$

$$= C + \frac{\frac{D_i z}{\omega_{a,t}}}{\frac{(1+r)(1-\delta_{a,t})-1}{(1+r)(1-\delta_{a,t})}} \quad (4.55)$$

$$= C + \frac{\frac{D_i z}{\omega_{a,t}}}{\frac{1+r-\delta_{a,t}-r\delta_{a,t}-1}{(1+r)(1-\delta_{a,t})}} \quad (4.56)$$

$$= C + \frac{(1+r)(1-\delta_{a,t})D_i z}{\omega_{a,t}(r - \delta_{a,t} - r\delta_{a,t})} \quad (4.57)$$

$$(4.58)$$

This converges if $|\frac{1}{(1+r)(1-\delta)}| < 1$, for which we require $r > \frac{\delta}{1-\delta}$.

4.G Optimal price derivation

The price for a car m is set to maximise expected profit in a given segment s :

$$\mathbb{E}(\Pi_{m,s,t}(P)) = \frac{(P - C_{m,t})|I_{s,t}| \exp(\kappa U_{m,s,t}(P))}{\exp(\kappa U_{m,s,t}(P)) + \sum_{m=0}^{Y_{t-1}} \exp(\kappa U_{m,s,t}(P))} \quad (4.59)$$

The optimal price satisfies profit maximization:

$$P_{m,s,t}^* = \operatorname{argmax}_P \mathbb{E}(\Pi_{m,s,t}(P)) \quad (4.60)$$

The price setter assumes the distance travelled equals the average distance of the entire population ($D_s = \sum_i^I D_i / |I|$) and disregards the price at which consumers can sell their car ($\hat{P}_{s,t} = 0$). In what follows we use the simplifying notation below:

$$\begin{aligned} O_{m,s,t} &= -\gamma_i E_{a,t} + \beta_i Q_{a,t}^\alpha + \nu (B_{a,t} \omega_{a,t} (1 - \delta_{a,t})^{L_{a,t}})^\zeta - D_i \frac{(1+r)(1-\delta_{a,t})(c_{a,t} + \gamma_i e_{a,t})}{\omega_{a,t}(r-\delta_{a,t}-r\delta_{a,t})} \\ W_{s,t} &= \sum_{m=0}^{Y_{t-1}} \exp(\kappa U_{m,s,t}) \end{aligned} \quad (4.61)$$

Substituting in the profit equation, we get:

$$\Pi(P)_{m,s,t} = \frac{(P - C_{m,t}) |I_{s,t}| \exp(\kappa(O_{m,s,t} - P))}{\exp(\kappa(O_{m,s,t} - P)) + W_{s,t}} \quad (4.62)$$

$$= \frac{(P - C_{m,t}) |I_{s,t}| h}{h + W_{s,t}} \quad (4.63)$$

where $h = \exp(\kappa(O_{m,s,t} - P))$,

$$\frac{\partial h}{\partial P} = -\kappa h \quad (4.64)$$

Take partial differential with respect to P and set to 0, solve for $P_{m,s,t}^*$.

$$\frac{\partial \Pi}{\partial P} = |I_{s,t}| \frac{h}{h + W_{s,t}} - (P - C_{m,t}) |I_{s,t}| \frac{\kappa W_{s,t} h}{(h + W_{s,t})^2} \quad (4.65)$$

$$0 = |I_{s,t}| \frac{h}{h + W_{s,t}} - (P_{m,s,t}^* - C_{m,t}) |I_{s,t}| \frac{\kappa W_{s,t} h}{(h + W_{s,t})^2} \quad (4.66)$$

$$0 = 1 - (P_{m,s,t}^* - C_{m,t}) \frac{\kappa W_{s,t}}{h + W_{s,t}} \quad (4.67)$$

$$h = W_{s,t}((P_{m,s,t}^* - C_{m,t})\kappa - 1) \quad (4.68)$$

$$\exp(\kappa(O_{m,s,t} - P_{m,s,t}^*)) = W_{s,t}((P_{m,s,t}^* - C_{m,t})\kappa - 1) \quad (4.69)$$

introduce $z = (P - C_{m,t})\kappa - 1$ where,

$$P_{m,s,t}^* = C_{m,t} + \frac{z + 1}{\kappa} \quad (4.70)$$

$$\exp(\kappa(O_{m,s,t} - P_{m,s,t}^*)) = z W_{s,t} \quad (4.71)$$

$$\exp(-\kappa(C_{m,t} + \frac{z + 1}{\kappa}) + \kappa O_{m,s,t}) = z W_{s,t} \quad (4.72)$$

$$\exp(-z - 1 - \kappa C_{m,t} + \kappa O_{m,s,t}) = z W_{s,t} \quad (4.73)$$

$$z \exp(z) = \frac{\exp(\kappa(O_{m,s,t} - C_{m,t}) - 1)}{W_{s,t}} \quad (4.74)$$

using the Lambert function \mathbb{W} ,

$$P_{s,m,t}^* = C_{m,t} + \frac{1}{\kappa} + \frac{\mathbb{W}\left(\frac{\exp(\kappa(O_{m,s,t} - C_{m,t}) - 1)}{W_{s,t}}\right)}{\kappa} \quad (4.75)$$

$$\begin{aligned} &= C_{m,t} + \frac{1}{\kappa} + \frac{\mathbb{W}\left(\frac{\exp\left(\kappa(-\gamma_i E_{a,t} + \beta_i Q_{a,t}^\alpha + \nu (B_{a,t} \omega_{a,t} (1 - \delta_{a,t})^{L_{a,t}})^\zeta - D_i \frac{(1+r)(1-\delta_{a,t})(c_{a,t} + \gamma_i e_{a,t})}{\omega_{a,t}(r-\delta_{a,t}-r\delta_{a,t})} - C_{m,t}) - 1\right)}{W_{s,t}}\right)}{\kappa} \end{aligned} \quad (4.76)$$

Thus the optimal profit function may be broken down into 3 components: the production cost, a bounded rationality aspect which tends to 0 as individuals become perfectly rational and a market advantage component which increases with car utility relative to the rest of the competition.

4.H Additional figures




Figure 4.H.1: *Histogram of willingness to pay for quality (top left), willingness to pay for emissions reduction (top right), EV imitation threshold (bottom left), distance driven (bottom right).*




Figure 4.H.2: *Posterior density for different values of α_χ, β_χ .*




Figure 4.H.3: *Monthly California Regular Reformulated Retail Gasoline Prices and Residential Electricity Price in Dollars per kilowatt-hour. Source: EIA.*




Figure 4.H.4: *Monthly California Regular Reformulated Retail Gasoline Prices and Residential Electricity Price in Dollars per kilowatt-hour. Source: EIA.*




Figure 4.H.5: *Value distribution in the NK model for different values of K -*




Figure 4.H.6: *EV uptake proportion for local variation of key parameters.*




Figure 4.H.7: *Sobol sensitivity using first order index.*




Figure 4.H.8: *Sobol sensitivity using total order index.*

Chapter 5

Conclusions

Addressing the climate crisis through a move towards low-carbon lifestyles requires not only personal changes but a systemic policy approach that induces collective transformation in behaviours. As this thesis has argued, achieving such a shift involves recognising the economy as a complex adaptive system, where behavioural change is shaped not only by short-term social influence but also by slower, longer-term cultural dynamics. To study these processes, this thesis applied agent-based modelling (ABM) to examine how individuals' pro-environmental attitudes, consumption habits, and firm innovation co-evolve under the influence of cultural dynamics and climate policies.

Chapter 2 examined how low-carbon consumption choices interact with cultural change, focusing on the evolution of individuals' pro-environmental attitudes and their influence on behavioural choices, such as whether to follow a vegetarian diet or cycle to work. These attitudes evolved due to information exchange via imperfect social learning in a social network of peers. The key model mechanism was that the intensity of these interactions was determined by the similarity over multiple pro-environmental attitudes, expressed collectively as an environmental identity. This led to inter-behavioural dependency and spillovers in green attitudes. Results showed that the initial distribution of agent attitudes towards behaviours and asymmetries in social learning were the main drivers of model dynamics, helping to generate awareness of what roadblocks may appear to widespread decarbonisation. Finally, the model incorporated green influencers, a minority of individuals who broadcast strong pro-environmental attitudes. This facilitated the study of how cultural dynamics can accelerate a green transition beyond the effects of social diffusion. In this scenario, including behavioural spillovers leads to greater emissions reductions than when behavioural choices are treated independently. Despite this, the inclusion of green influencers failed to produce deep behavioural decarbonisation through solely voluntary action. This indicated that information provision policies alone are insufficient and must be complemented by policies that reduce barriers to adopting pro-environmental behaviour.

A key takeaway from Chapter 2 is the limited decarbonisation potential of relying solely on voluntary behavioural change through information diffusion. These results highlight the importance of tailoring interventions to individuals' environmental identities to avoid disengaging those who might rebuke pro-environmental information if it is too green. Additionally, the findings suggest that information provision policies must be embedded within a supportive policy environment that lowers structural barriers to low-carbon choices, such as carbon pricing or industry standards. Without this, social influence alone is unlikely to generate the scale of behavioural change needed for meaningful

emissions reduction.

Chapter 3 was an extension of Chapter 2, keeping the same core model but using a nested constant elasticity of substitution utility function to determine the quantities of low- or high-carbon goods consumed. This adjustment made it possible to study how cultural change affects consumption when climate policies like carbon taxes are introduced. The notion of a “social multiplier” of environmental policy was extended to the context of multiple consumer needs while allowing for behavioural spillovers between these, giving rise to a “cultural multiplier”. Results showed that the cultural multiplier stimulates greater reduction in emissions compared to the case with fixed preferences. Additionally, it was found that the cultural multiplier greatly increases the effectiveness of a relatively small carbon price. This effect may improve the acceptability of carbon pricing by reducing the level of stringency required to achieve the same decarbonisation target. At high carbon tax levels, the distinction between social and cultural multiplier effects became less pronounced. The strong price signal, amplified through social imitation, was sufficient to drive behavioural change regardless of the presence or absence of behavioural spillovers. As a result, even individuals with low pro-environmental preferences adopted low-carbon consumption. Moreover, by varying socio-economic conditions, such as substitutability between low- and high-carbon goods, social network structure, proximity of like-minded individuals and the diversity of consumption lifestyles, the model provides insights into how cultural change can be leveraged to induce maximum effectiveness of climate policy.

The findings of Chapter 3 highlight the importance of complementing carbon taxes with additional policy instruments that foster stronger pro-environmental identities to increase decarbonisation potential. This may take the form of expanding current visions of what individuals consider to be key consumption behaviours for a low-carbon lifestyle. For example, information provision policies such as eco-labelling may correct misinformation on the true carbon impact of less socially salient consumption categories. Furthermore, increasing the substitutability between low- and high-carbon goods, either via technological improvements or behavioural nudges to increase the appeal of low-carbon alternatives, can also enhance the cultural multiplier. Lastly, high similarity in pro- or anti-environmental identities within communities can hinder decarbonization efforts. This makes it important for policymakers to consider the network structures that influence social imitation when evaluating the potential impact of carbon taxation.

Chapter 4 extended the adaptive complexity perspective to the co-evolution of technological innovation by firms and consumer behaviour. As in Chapters 2 and 3 it maintained a focus on the role of social learning in individual decision-making. The chapter developed a novel ABM of the diffusion of electric vehicles (EVs) in California calibrated on the period 2001-2023. In the model, heterogeneous individuals influenced by their social peers’ purchase cars, both new and used, while manufacturers develop new car models through the exploration of an NK landscape. Different policy combinations were considered when striving to achieve ambitious EV deployment targets and simultaneously balancing economic costs, emissions and consumer utility. Simulations compared individual and combined policy instruments, including carbon pricing, new and used car purchase rebates, production subsidies and electricity price subsidies. Results indicate that only carbon pricing and new vehicle purchase rebates as single policy instruments can achieve a target of 95% electric vehicle adoption by 2035. Each policy has key trade-offs: carbon pricing substantially reduces consumer welfare, while new purchase rebates create heavy fiscal burdens. Policy combinations prove to be more effective by reducing the stringency of the required interventions and minimising the induced rebound effect

and trade-offs. Notably, policies that are ineffective in isolation, such as production and electricity subsidies, become viable when implemented in a policy mix. Extended simulations till 2050 demonstrated that successful transitions to an EV fleet stabilise after policy withdrawal. Yet this rapid transition comes with a trade-off that the cumulative emissions in the business-as-usual scenario are lower than those of a mobility transition until the 2040s due to high production emissions for new EVs.

Chapter 4 demonstrates that for a rapid transition to an EV fleet highly stringent policies are required. Only by harnessing synergies between policy instruments can negative trade-offs in cost and welfare be minimized. Additionally, the findings highlight that an accelerating EV adoption alone is insufficient for producing sustainable mobility, due to persistent production-related emissions and the structural dependence on private car use. This calls for a broader policy agenda that targets the root of the problem by reducing car dependence altogether. This might be achieved through improved public transport, urban redesign, and alternative mobility models such as car sharing.

Taken together, the three chapters underscore the limits of isolated policy instruments in achieving widespread decarbonisation. By recognising the economy as a complex adaptive system, this thesis argues for a more holistic approach to climate policy. Therefore, the analysis of climate policy must also move beyond static models. Instead, future models should account for the co-evolution of consumer and firm behaviour by embedding processes of social learning, cultural dynamics and technological innovation.

Overall, the findings of this PhD thesis contribute new ABMs and policy insights to the literature, emphasizing the importance of complementing top-down approaches to climate policy analysis with bottom-up perspectives. This is due to the richer representation of heterogeneity and the complexity of social interactions that bottom-up approaches can provide. Whilst the ABMs described improve on certain aspects of top-down models, they are still susceptible to several key flaws common to both approaches. First, there is a tendency to centre those dynamics which may be easily quantified or represented in equation form. For example, Chapter 4 omits non-market-based policies such as such as car sharing or remote working which despite their potential to reduce mobility demand, are challenging to incorporate within the ABM framework. Second, in the pursuit of greater model realism, there can be an explosion in complexity. This easily leads to opacity in insights resulting in the ABM becoming a black box. Additionally, high model complexity results in a reduced ability to perform global sensitivity analysis due to high computational cost. Third, complex models can be difficult to validate when they have a very large number of parameters, which are hard to calibrate from the literature. This highlights the need to strike a balance between realism and model opacity. Models should be complex enough to capture key dynamics, avoiding self-evident conclusions, but not too complex that their results become uninterpretable.

The model in Chapter 2 adopts several simplifications to manage complexity, which future research could address to improve behavioural realism. For instance, many pro-environmental behaviours (such as purchasing an EV, installing solar panels, or cycling to work) are discrete, independent, and often path-dependent. While the model's current continuous representation offers a useful starting point, future models could incorporate more detailed structures to reflect these behavioural characteristics better.

As a future extension to the model in Chapter 3, the inclusion of rebound effects would better capture the complexity of real-world consumption dynamics. Additionally, using a utility function with non-homothetic preferences would allow for a representation of the differing capabilities of households across income deciles to shift towards low-carbon

options. This extension would give valuable insight into how economic inequality interacts with cultural dynamics in shaping behavioural change.

The next step in future research for Chapter 4 would be integrating forward-looking behaviour through expectation formation in vehicle users and firms. This would enable the assessment of command and control policy instruments such as emissions standards or zero-emission vehicle mandates, as well as allow for gradual increases in policy intensities. Additionally, incorporating supply-side frictions such as capital constraints and new car inventories would also enhance the realism of the transition path, particularly in reflecting bottlenecks that may emerge in scaling up EV manufacturing. Lastly, future work should incorporate alternative modes of mobility such as cycling, public transport or car sharing to model the move away from car reliance.