



Politecnico
di Torino

ScuDo

Scuola di Dottorato - Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Program in Management Engineering (38th cycle)

Integrating complexity in innovation diffusion processes

By

Giacomo Masali

Supervisor(s):

Prof. Guido Perboli, Supervisor

Prof. Maria Elena Bruni, Co-Supervisor

Doctoral Examination Committee:

Prof. François Lafond, Referee, INET, Oxford University

Prof. Sara Khodaparasti, Referee, Università Mercatorum

Politecnico di Torino

2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Giacomo Masali
2026

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

A nonno Luciano

Acknowledgements

I would like to thank my supervisor professor Guido Perboli and co-supervisor professor Maria Elena Bruni for the constant help and assistance that they provided throughout the three years of the PhD program. I thank professor François Lafond and doctors Joel Dyer and Jordan Kemp for the insightful discussions and feedback that shaped the direction of this work. I am also grateful to professor Doyne Farmer and to all the people of the Complexity Economics group at INET Oxford for the helpful comments during group meetings and for the stimulating research environment that I have found during my visiting period. Any remaining errors are my own.

Abstract

The diffusion of innovation is a paradigmatic complex-systems phenomenon: aggregate S-curves emerge from the interaction of heterogeneous agents, evolving technologies, institutions, and media. This thesis uses tools from complexity science to connect three levels of description that are usually treated separately: parsimonious aggregate diffusion curves, micro-founded kinetic models, and empirically grounded agent-based simulations. The first part takes a deliberately reduced-form perspective. Building on the diffusion literature in operations research and marketing, it fits a Bass-type model to SDG Index trajectories (2000-2021) for each country, treating the parameters as mechanics-agnostic summaries of aggregate dynamics. Diagnostics such as saturation time and the timing of peak annual improvement reveal that most trajectories follow the familiar pattern of slow start, acceleration, and saturation, with systematic differences across income groups and regions. This provides a quantitative baseline that complements the largely qualitative work on national innovation ecosystems and sustainability. The second part moves from description to mechanism via kinetic theory. Inspired by the kinetic modeling of socio-economic systems, it develops a Kinetic Innovation Diffusion (KID) model in which individuals carry a perceived value for the product and a binary adoption state. Social interaction, media exposure, and experience are modeled as “collisions” that update perceived values; in a suitable limit, the Boltzmann description yields a closed-form adoption law where the S-curve arises as a regularized incomplete gamma function with explicit dependence on media, performance, and heterogeneity. Calibrated on four benchmark consumer durables, the KID model matches or outperforms standard Bass-family models and supports an early-stage forecasting scheme that combines pre-launch belief data with a small number of sales observations. The third part extends this kinetic framework to expectations and technological performance, providing a micro-founded model of the Gartner Hype Cycle, whose original formulation remains largely heuristic and empirically

debated. Mean perceived value among non-users (expectations), adoption, and realized performance are coupled in a dynamical system: media and word-of-mouth raise expectations; adoption exposes users to realized performance; if performance lags expectations, negative feedback generates a crash and subsequent recovery. In a deterministic limit, this structure yields transparent conditions on the relative growth of media, adoption, and performance for the peak and the trough; reintroducing heterogeneity restores smooth S-shaped adoption. A Direct Simulation Monte Carlo implementation shows that hype-disappointment-recovery patterns persist without artificial time-scale separation and even under constant media, turning the Hype Cycle from a descriptive curve into an emergent behavior with falsifiable implications. Finally, the thesis grounds these mechanisms in an agent-based model of Connected, Communicative, Automated Mobility diffusion in Europe, using a synthetic population derived from focus-group data on socio-demographics, mobility habits, and attitudes. Agents have mobility priorities and a sentiment toward CCAM; the service is characterized by performance and communication parameters. The model jointly captures mode choice and adoption/abandonment of CCAM, with sentiment evolving through experience and word-of-mouth, echoing the hype dynamics analyzed earlier. Scenario analysis shows that increasing CCAM speed alone yields progressively more successful diffusion but limited and fragile mode shift away from cars; to meaningfully reduce car use, CCAM must also improve reachability and comfort, which are dominant priorities for car users. More broadly, the ABM acts as a scenario laboratory where product design and communication strategies can be stress-tested against heterogeneous preferences and expectations. Overall, the thesis advances a unified view: diffusion of innovation is best modeled as a complex system. Aggregate S-curves, kinetic equations, and agent-based simulations are not competing paradigms but complementary lenses on the same process, enabling richer description, mechanism-based explanation, and policy-relevant experimentation under uncertainty.

Contents

List of Figures	x
List of Tables	xii
1 Introduction	1
1.1 Diffusion of innovations	1
1.2 The science of complex systems	3
1.3 The diffusion of innovation is a complex system	5
1.4 Structure of the thesis	6
2 The Bass model for SDG Saturation	8
2.1 General Context and Related Literature	8
2.2 Data and Methodology	12
2.2.1 Data	12
2.2.2 The Bass Diffusion Model	13
2.2.3 The Adaptation of the Bass Diffusion Model	15
2.3 Results	16
2.4 Discussion	18
2.5 Conclusions	20
2.5.1 Limitations and Future Research Directions	23

3	The statistical mechanics of innovation diffusion	24
3.1	General Context	25
3.2	Related Literature	27
3.3	Kinetic Framework	29
3.3.1	The Boltzmann equation	29
3.3.2	Binary interactions	30
3.3.3	Label switch process	32
3.3.4	Time-scale separation	34
3.3.5	Quasi-invariant limit	35
3.4	Multi-agent Bass model	36
3.5	Results	44
3.5.1	Nonlinear Least Squares	44
3.5.2	Bayesian Inference	45
3.5.3	Early Stage Forecasting Method	48
3.6	Conclusions	51
4	The Hype Cycle as an Emergent Behavior	53
4.1	General Context	53
4.2	The Hype Cycle Framework	55
4.3	Model	58
4.4	Analytical Results and Simulations	62
4.4.1	Deterministic case	62
4.4.2	Noisy case	65
4.4.3	Agent-based simulation	67
4.4.4	Figures Generation	67
4.5	Conclusions	69

5	An Agent-based Model for the Diffusion of Automated Vehicles	72
5.1	General Context	72
5.2	Methodology	74
5.2.1	Data	74
5.2.2	Agent-based model	76
5.3	Results	80
5.3.1	Scenario 1	81
5.3.2	Scenario 2	83
5.3.3	Scenario Generation	84
5.4	Conclusions	85
6	Conclusions	87
6.1	Main Contributions	87
6.2	Limitations	89
6.3	Future Developments and Final Remarks	90
	References	93
	Appendix A SDG Index Peak and Saturation	106
	Appendix B Posterior Distributions	108

List of Figures

2.1	SDG Index Score progression and fitted Bass curve for four countries between years 2000 and 2021.	17
2.2	SDG Index Score progression and fitted Bass curve for four aggregates between years 2000 and 2021.	18
3.1	Example of perceived value distribution in the population.	41
3.2	Possible diffusion S-curves obtained by Bayesian inference of the parameters on the available data points (in black).	49
4.1	The Hype Cycle formed by the hype curve and the technological S-curve and the five phases. Adapted from Fenn and Time (2007).	56
4.2	The evolution of expectations, adoption and realized performance in the dynamical system.	65
4.3	The evolution in time of expectations (blue), adoption (red) and performance (green) obtained with the DSMC.	71
5.1	BIC score for different numbers of clusters.	75
5.2	The attributes of the survey data (first and third lines) and the corresponding attributes of the agents (second and fourth lines).	76
5.3	Evolution of the diffusion process with number of users (<i>non users</i> in blue, <i>users</i> in red and <i>haters</i> in yellow) and travel mode usage (car users in red, bike users in orange, public transport users in blue and walkers in yellow), for different values of Speed of the CCAM system.	82

5.4	The evolution of the average sentiment following the Hype Cycle pattern.	84
A.1	Legend.	106
A.2	The fitting parameters, peak and saturation times and fitting error for the first third of the analyzed entities. The RMSEs in yellow are between 0.15 and 0.2, the ones in red exceed 0.2.	106
A.3	The fitting parameters, peak and saturation times and fitting error for the second third of the analyzed entities. The RMSEs in yellow are between 0.15 and 0.2, the ones in red exceed 0.2.	107
A.4	The fitting parameters, peak and saturation times and fitting error for the last third of the analyzed entities. The RMSEs in yellow are between 0.15 and 0.2, the ones in red exceed 0.2.	107
B.1	The posterior distributions of the parameters obtained through the Metropolis-Hastings algorithm, for the products color TV and air conditioner.	108
B.2	The posterior distributions of the parameters obtained through the Metropolis-Hastings algorithm, for the products dryer and freezer.	109

List of Tables

2.1	Estimated Bass parameters, relative peak and saturation years and RMSE for the countries and aggregates in figures 2.1 and 2.2	19
3.1	The SSE for the four products of the compared models.	45
3.2	The parameters of the kinetic diffusion model for the four products.	45
3.3	The Sobol analysis results for each product.	47
3.4	Mean and standard deviation of forecast values of market potential and market saturation with 5 and 3 years of data compared with the values evaluated with the full sample, for the four products.	50
3.5	Mean and standard deviation of forecast values of sales peak and flex point with 5 and 3 years of data compared with the actual values, for the four products.	51
4.1	Perceived value updates ($v \mapsto v'$) by group and channel.	59
5.1	Satisfaction thresholds for mode-priority pairs.	77
5.2	CCAM system parameters for different simulations.	81

Chapter 1

Introduction

In this dissertation we use the lens of complexity to look at how new ideas and technologies spread. Rather than treating innovation diffusion and adoption as a simple curve on a chart, we follow the many small interactions that add up to visible change. First, we show when a compact description well captures the shape of innovative processes. Then, we open the process up, tracing how expectations and experience interact, and building population-level scenarios to see how outcomes change when conditions change. The goal is a coherent, readable account of why the diffusion of innovation takes the paths it does, and how that understanding can be used fruitfully.

1.1 Diffusion of innovations

How new things are adopted by people and spread in a population is a complex and fascinating phenomenon that is studied in a broad and diverse literature that spans different fields. Its earliest roots are sociological: Gabriel Tarde argued that social change propagates via *imitation*, the micro-to-macro process through which practices and ideas spread across a population that today we call *adoption*. His *Laws of Imitation* (Tarde, 1903) framed diffusion as a patterned, contagious behavior rather than isolated decisions and paved the way to the scientific study of the diffusion of innovations. He went as far as recognizing the typical S-shaped curve that describes the number of adopters of an innovation in time and to observe that the takeoff occurs when the opinion leaders in the system adopt, foreshadowing the relevance

that network theory has today in diffusion studies. Albeit seminal, Tarde's ideas could not rely on the quantitative methods that support modern diffusion studies, and remained marginal for half a century.

It was the rural sociology tradition that, in the 1940s, re-ignited the interest in the study of the diffusion of innovations with a seminal study on hybrid corn diffusion in Iowa, United States. Working in two farming communities, Ryan and Gross (1943) reconstructed who first heard about hybrid seed corn and from whom, and when farmers moved from trialing a few rows to planting most of a field. What emerges is not a tale of instantaneous conversion but a decade-long relay in which communication channels shift over time: early awareness often arrived via seed salesmen and extension agents, but adoption decisions crystallized through conversations with known adopters: neighbors whose credibility is earned by harvests rather than the promises of salespersons and media campaigns. The aggregate curve they documented is again the S-curve: a long, hesitant beginning, a sharp middle acceleration as interpersonal influence compounds, and a late saturation as the remaining holdouts convert. By placing interpersonal ties at the center, without denying a role for mass media, the study demonstrated that diffusion is a social process unfolding through social contact, and it seeded the field's canonical ideas about adoption stages and categories that Rogers (2003) later systematized. This is the moment when diffusion research became a coherent paradigm rather than scattered anecdotes.

Parallel to the sociological tradition, marketing science developed aggregate, model-based descriptions of sales and adoption. The touchstone is Bass (1969), which explains the typical S-curve as the interaction of two forces: an *innovation* effect (external influence) and an *imitation* effect (internal/social influence). The Bass model is a simple and parsimonious model that, using just three parameters, fits a wide range of product classes and remains a baseline for forecasting peak timing and market saturation. Over time, the Bass family expanded to handle realities of modern markets: marketing-mix covariates, competition and network externalities, and multi-generation product lines.

Since the 1990s, diffusion research has broadened from "word-of-mouth on a homogeneous market" to a networked, multi-market, multi-generation view. Peres et al. (2010) synthesize this shift. Inside a market, diffusion studies now foreground social networks, network externalities, takeoffs and saddles, and technology gener-

ations. Agent-based models (ABMs, for short) and choice–diffusion hybrids (e.g., Markov chains) connect individual rules to aggregate growth and separate interpersonal influence from externalities; saddles and chasms emerge naturally from heterogeneity and interactions rather than from noise. Across markets, research models cross-country lead-lag effects, systematic growth differences by culture or economy, and competition’s imprint on category and brand growth. The common thread is a re-framing of diffusion as social interdependence of all kinds, studied with network data, ABMs, and multinational models.

In sum, the diffusion research gives us two complementary insights that cut across all fields: first, there is rich evidence that adoption is fundamentally social and context-dependent; and second, parsimonious aggregate models are extremely useful to summarize and forecast diffusion, but they do not expose the complexity of the underlying social mechanisms. This is precisely why later work often pairs aggregate models with micro- or meso-level mechanisms such as agent-based models, network theory, or, as we will see in chapter 3, kinetic theory. This is where the science of complex systems comes into play: it offers the methods to keep those mechanisms explicit while remaining analytically disciplined and rigorous. It is through the lenses of complexity that diffusion of innovation appears as an emergent property of a far richer system.

1.2 The science of complex systems

Complex systems science begins from the observation that when many interacting parts update in response to one another, the collective acquires properties one cannot deduce by inspecting the parts in isolation. These emergent behaviors – phase changes, self-organization, persistent path-dependence – arise because local interactions and feedback loops amplify small differences and create nonlinear cause-and-effect: a tiny perturbation can trigger a qualitative shift, while large shocks may be absorbed with little visible change. The field’s original purpose was exactly this unification: to provide a common, predictive language for interacting systems on the premise that similar patterns recur across domains once heterogeneity, feedback, and topology are modeled explicitly. The aim is not a grand theory of everything, but a disciplined framework that (i) specifies how elements influence one another, (ii)

derives the emergent macrostates that follow, and (iii) connects those macrostates to testable predictions.

The logic of complexity recurs across very different arenas. In biology, local biochemical interactions and cell to cell signaling generate organism-level form and function: from morphogenesis to neural coordination, macroscopic order (organs, behaviors) arises without a central designer, via self-organization and cascading feedbacks that amplify tiny molecular asymmetries into qualitative shifts in phenotype. In epidemiology, micro-level contacts and behavioral responses produce epidemic curves whose timing and height depend nonlinearly on small changes in exposure, precaution, or environment. In financial markets, heterogeneous traders reacting to each other's moves drive volatility clustering and regime switches: modest informational shocks can either dissipate or snowball into large swings depending on state and feedback. What ties these domains together is the same structure (local interactions, feedback, and nonlinearity), which explains why similar phenomena (S-shaped build-ups, oscillations, abrupt transitions) appear despite different constituents and domains.

Methodologically, the study of complex systems lives at the intersection of mathematics and computation. On the mathematical side, network theory gives a compact way to say who interacts with whom and how influence accumulates; statistical mechanics translate many small interactions into mathematically tractable laws for the whole, where non-linear and dynamical systems analysis often reveals thresholds, phase changes, and long-lived regimes. On the computational side, modern power and data have shifted the center of gravity: agent-based models (ABMs) let us write simple local rules for heterogeneous agents and watch macro patterns emerge; Monte Carlo methods navigate nonconvex fitting problems, solve non-linear differential equations and quantify uncertainty; large datasets make calibration and out-of-sample testing of large models possible. Together, these methods let us move smoothly between micro rules and macro regularities, preserving the essence of complexity – the way wholes acquire properties irreducible to their parts – without sacrificing predictive discipline.

To conclude, complex systems science gives us a way to think clearly about phenomena whose visible regularities are born from invisible interactions. It explains why smooth aggregate curves can hide sharp local shifts, why timing matters as much as magnitude, and why the same few ingredients reappear across fields. With

that lens in hand, it is natural to treat diffusion of innovation as a complex process in which expectations, experience, and behavior coevolve.

1.3 The diffusion of innovation is a complex system

If you watch a market from close enough, it refuses to sit still. Consumers compare experiences and share information, firms improve products, journalists and influencers broadcast partial truths, prices move, and yesterday's decision becomes today's prior. The moving parts are many and none of them act in isolation. Signals provoke behavior, behavior feeds back into signals, and network structure quietly shapes what becomes visible. What we call the diffusion of an innovation is the visible trace of all this hidden motion. This is the natural habitat of complex systems. In markets, a single enthusiastic review can trigger a local cascade; a disappointing early experience can poison a cluster for months. The same advertising budget can accelerate adoption in one community and backfire in another, depending on who is connected to whom and how credibility flows. Complexity is not a flourish added to diffusion; it is the ambient condition under which diffusion takes place.

What follows in the thesis is a progressively sharper look at diffusion-as-a-complex-system, where we employ a combination of analytical and computational tools typical of the science of complex systems. We exploit the parsimonious macro descriptors of the Bass model to show that, even in messy, multi-actor environments such as national innovation ecosystems (NIE), simple S-curves can summarize trajectories and forecast turning points. This is an innovative field of application of the Bass model, that still yields good results. Crucially, this exposes the fact that such top-down aggregated models are mechanics-agnostic and, albeit useful, do not provide insights on the microscopic mechanisms of a rich, non-linear process. Consequently, we employ methods from statistical mechanics to model the diffusion of an innovation from the bottom up. Starting from the microscopic interactions between boundedly rational agents in a population, we use the kinetic theory to draw a closed-form equation for the adoption of a new product in a market. To show the good descriptive power of the model, we test it against state of the art Bass-type models on four benchmark products. The calibration of the parameters is performed using a Markov Chain Monte Carlo algorithm for Bayesian inference. The same model can be expanded to account for other social and technological

processes related to diffusion: the dynamics of people expectations towards an innovation and of technology performance development are added to the picture. The resulting dynamical system provides a rigorous mathematical explanation for an influential observed pattern in innovation studies, the Hype Cycle, which lacked such a formalization. A Direct Simulation Monte Carlo of such a model is also employed as validation of the analytical results and to investigate peculiar parameter combinations. Additionally, to fully account for heterogeneity and the use of large datasets, we present an example of agent-based model for the diffusion of automated mobility in Europe.

In sum, diffusion is a complex system: boundedly rational heterogeneous agents, interactions, and feedback generate patterns no top-down curve can explain alone. We use different tools for different scales – S-curves, kinetic theory, Monte Carlo computation, dynamical systems and ABM – to make those mechanisms visible and testable.

1.4 Structure of the thesis

This dissertation moves from macro description to micro-founded mechanism and is organized in thematic chapters. In each chapter, the relevant literature is presented contextually to the topic of the discussion, the methodology is thoroughly explained, and the contributions and results for each chapter are clearly stated. The structure of the thesis is the following:

- **Chapter 2** introduces the classic Bass framework and applies it to the forecasting of levels and times of fulfillment of the Sustainable Development Goals (SDGs) by different nations. It is shown how, apart from a few specific cases and after pre-processing and smoothing of the raw data, the S-curve typical of the Bass model well describes the progress in time of the global SDG score for a majority of countries. The point is not to claim mechanism, but to extract stable diagnostics such as takeoff, peak improvement, and saturation that policy makers can read.
- **Chapter 3** takes an inverted perspective with respect to the Bass model. The process of diffusion is modeled from the bottom up to expose the microscopic mechanics that let the S-curve emerge. The kinetic theory is introduced

and used to obtain a differential equation for the adoption in time. Then, computational results of parameter estimation via Bayesian inference and forecasting performance are shown.

- **Chapter 4** expands the model of the previous chapter to the dynamic of expectations in the population, to provide a mathematical model for the formation of Hype Cycles. First, the concept of Hype Cycle is introduced and then it is explained as an emergent behavior produced by the interplay of expectations, adoption and technology performance.
- **Chapter 5** constitutes an example of agent-based model of diffusion of innovation. Data on consumer and transportation habits from SINFONICA, the flagship European-funded project that aims to foster the diffusion of Connected, Communicating and Automated Mobility (CCAM) systems in Europe, are used to build a synthetic population of agents that then is left to interact and take action following a set of simple rules. Two scenarios, with different performance levels of the CCAM service that produce qualitatively different diffusion patterns are discussed.
- **Chapter 6** concludes the dissertation delineating the perimeters of the research, synthesizing the original contributions, and indicating the directions for future research.

Chapter 2

The Bass model for SDG Saturation

In this chapter, we show how a top-down, aggregated model such as the Bass model can be successfully employed to capture the visible trace of a complex system different from the one that was originally intended. We use the Bass S-curve not to explain mechanisms, but to summarize how national innovation ecosystems (NIEs) advance toward the Sustainable Development Goals (SDGs). The claim is modest and useful: if many heterogeneous actors, feedbacks, and policies collectively produce an S-shaped trajectory at the country level, a parsimonious curve can still forecast turning points and saturation, even though the underlying machinery differs from consumer adoption. The content of this chapter is taken from Masali et al. (2026).

2.1 General Context and Related Literature

A national innovation ecosystem (NIE) is best understood as an open, adaptive arrangement of legally independent actors such as public agencies, firms, universities, financiers, and civil society, whose ongoing collaboration and exchange of knowledge generate and scale innovations. Rather than a top-down hierarchy or a static system, an NIE behaves like a complex adaptive system: it self-organizes, adapts to shocks, and exhibits emergent properties (synergy, path dependence) because interactions and feedbacks among participants continuously reshape capabilities and incentives. In this view, innovation is co-created through complementary assets, shared infrastructures, and relational contracts and not simply produced by isolated champions. Policy shifts worldwide reflect this turn from supporting individual

agents to orchestrating collaborative entities (clusters, platforms, multi-stakeholder partnerships) that link localized networks into economy-wide ecosystems. Formation is thus bottom-up and top-down at once: local collaborative networks (e.g., clusters) accumulate and interlink; national strategy, institutions, and enabling rules lower coordination costs and stabilize cooperation; together they yield a multi-level ecosystem that evolves over time.

Mechanistically, NIEs work through repeated interaction and complementarities: shared standards and data, co-specialized investments, joint problem-solving, and iterative learning. These features generate nonlinearity (small coordination wins can unlock large diffusion effects) and make governance context-dependent: outcomes hinge on how well incentives align, how credible partners are, and how flexibly institutions support collaboration.

The Sustainable Development Goals are the UN's universal agenda for 2015–2030: 17 goals with 169 targets, designed to be integrated and indivisible across three pillars (economic, social, and environmental) while advancing a broader vision often summarized as people, planet, prosperity, peace, and partnership. They build on the Millennium Development Goals, pairing outcomes (e.g., poverty, health, education, climate) with means of implementation and a revitalized Global Partnership to mobilize finance, technology, and policy coherence. Countries commit to nationally led strategies and shared monitoring to ensure progress.

NIEs can accelerate technological, social, and economic progress toward achieving SDGs by fostering collaboration, innovation, and entrepreneurship (Azmat et al., 2023; Brás and Robaina, 2024; Iizuka and Hane, 2021; Liao et al., 2024; Liu and Stephens, 2019; Nylund et al., 2021). Specifically, a study on 131 countries demonstrates that NIEs improve the achievement of SDGs (Wei et al., 2025) through the development of eco-innovations and green technologies (Abdullahi et al., 2022; Yikun et al., 2023) entrenched with the tackling of societal issues and grand challenges (Ritala, 2024). Nevertheless, their effectiveness in regard to SDGs can be constrained by various barriers as well as complexity issues, including contingent approaches in innovation-as-a-context ecosystems, institutional inertia, resource limitations, and inadequate and nonlinearly intertwined efforts among the multiple agents in NIEs (Gifford et al., 2021; Grama-Vigouroux et al., 2023; Huang et al., 2024; Ponsiglione et al., 2021; Ponta et al., 2023; Qiao et al., 2024; Thomas et al., 2025). For instance, high-income countries show 2.3 times greater SDG progress per dollar invested in

innovation activities than low-income countries, due to better institutional frameworks that mitigate the complexity of NIEs (Prokop et al., 2021). This trend is confirmed by the overall dominance of the agenda for implementing circular ecosystems in the Global North (Aryee et al., 2025). The same research indicates that southern/emerging economies are both under-studied and largely necessity-driven arenas for circular innovation ecosystems, and that implementation tends to be more turbulent. The key differences concern how these ecosystems emerge, are orchestrated, and are governed. Telling examples relate to sectors such as organic food in Brazil (Ferrari et al., 2023), textiles and apparel in Romania (Staicu and Pop, 2018), biochemical energy in India (Mukherjee et al., 2021), and construction in Ghana (Gyimah et al., 2025). Evidence from complex contexts—such as Brazil and Pakistan—shows that participation in NIEs can lift SME productivity by 22–40% and cut resource waste by 15% (Jan et al., 2025). Additionally, the World Economic Forum links weak emphasis on sustainability-oriented innovation to deteriorating SDG performance (World Economic Forum, 2025).

Conversely, NIEs are most effective at driving sustainability when they embed inclusive, structured governance that coordinates multiple actors within complex systems (Wei et al., 2025; World Economic Forum, 2020). Over the past decade, work on innovation ecosystems has increasingly adopted a complexity lens, treating them as evolving configurations of actors, activities, and relationships, and has asked how structure and dynamics translate into outcomes. The emphasis has been on mapping who does what with whom, how those interactions change over time, and how these patterns drive performance and growth (Ancona et al., 2023; Marinelli et al., 2023; Neto et al., 2024). Moreover, several definitory efforts have been undertaken in order to identify the differentiating elements of innovation ecosystems, such as actors' heterogeneity, actors' interdependence, and system-level outcomes, all of them calling for a specific focus on complexity issues (Thomas et al., 2025). Innovation ecosystems rest on coevolutionary microfoundations: repeated interactions among firms, institutions, and other actors generate nonlinear feedbacks that produce emergent outcomes (Breslin et al., 2021). The resulting system-level patterns vary by ecosystem type – innovation, business, platform/technology, entrepreneurial, or knowledge and open-innovation – because each type organizes complementarities and coordination differently. In the same spirit, Reed et al. (2025) reject linear pipeline views, arguing for approaches that foreground policy shocks and adaptive strategies as triggers of cascading transformations across interconnected ecosystems.

Zhang et al. (2024) prove that green innovation spreads through unequal connection nodes in complex networks, arising nonlinearly from actors' micro-interactions (e.g., firms, universities, policymakers). Such interactions may include different agents' activities, such as pollution taxes or default costs altering opportunity costs, thus incentivizing green cooperation in NIEs, policy shifts, or startup disruptions. They can also set off disproportionate effects across the ecosystem, catalyzing learning and co-evolution among firms, governments, and universities that yields emergent behaviors that challenge linear forecasts. In parallel, growing emphasis on the sustainability transition in policy and corporate strategy fosters self-organization around shared goals, enabling organic reconfiguration of roles and complementarities (Han et al., 2021).

Recent work converges on a clear view: innovation ecosystems and sustainability are inseparable, and effective practice requires adaptive strategies that accommodate organizational and ecosystem complexity (Pham and Vu, 2022; Ritala, 2024). Sustainability rises on the research agenda because collaborative innovation generates environmental and social externalities alongside benefits (Liu and Stephens, 2019; Nylund et al., 2021), demands cross-sector coordination such as industrial symbiosis (Liu et al., 2025), and must navigate shifting policies and markets (Agasty et al., 2023). It also calls for embedding sustainability in business models and regional strategies (Oskam et al., 2021; Peñarroya-Farell et al., 2023). Yet the literature remains fragmented: governance and operating logics are under-specified (Thomas et al., 2025; Xu and Li, 2025), and a coherent framework for nonlinear, co-evolutionary interactions and their triple-bottom-line consequences is still emerging. A parallel strand highlights a dark side: some innovation pathways can impede SDG progress (Berry et al., 2024; Islam, 2025). Empirically, findings differ by context (high- vs. low-income, stable vs. post-conflict) producing mixed assessments of how NIE complexity affects SDG attainment (Grama-Vigouroux et al., 2023; Jan et al., 2025; Prokop et al., 2021). Consensus is narrow but meaningful: collaboration enables sustainability, while outcomes hinge on context-specific governance and ecosystem complementarities that shape system-level value creation.

A practical gap motivates what follows: despite the surge of interest in innovation ecosystems, quantitative methods remain underused. Much of the field still privileges conceptual mapping and narrative reviews over empirical testing and prediction, leaving open how far data-driven models can organize evidence and forecast trajectories. In this chapter, we take a deliberately reduced-form approach

and ask whether a simple aggregate diffusion law can capture the visible trace of national progress on the SDGs, without pretending to identify the full micro-causal machinery beneath it. Concretely, we assemble country-level SDG Index trajectories (2000–2021) and fit a Bass-type S-curve to each series, treating the parameters as summaries of aggregate dynamics. The central diagnostic is the saturation time – when progress approaches a country-specific ceiling – alongside the timing of the maximum year-over-year improvement. The value of this exercise is comparability: a common curve, estimated in a disciplined way, yields a baseline that travels across countries and makes planning and prioritization more coherent, even when underlying mechanisms differ. The preview of results is straightforward. For most countries, SDG progress displays the familiar sequence of slow start, acceleration, and saturation; a small subset (13) lacks a clear S-shape. Cross-country differences cluster along geographic and socioeconomic lines: places with stronger innovation capacity and policy effort tend to reach acceleration earlier, while those already near the ceiling exhibit lower remaining potential despite high inputs. In short, a mechanics-agnostic S-curve offers a useful, quantitative baseline for national SDG progress, precisely the kind of empirical scaffold the literature has been missing.

2.2 Data and Methodology

2.2.1 Data

The dataset analyzed for this research is the online database for the Sustainable Development Report 2022 by Sachs et al. (2022), which is a yearly report that reviews the progress made periodically by UN member states with regard to the level of SDG achievement since their adoption. The data used for this study span 22 years, from 2000 to 2021, and the dataset contains 120 indicators covering the 17 SDGs, which generate the overall score for each country. Overall, 193 countries are included in the Sustainable Development Report 2022. The dataset includes several data on the SDGs, such as the overall results for all countries, including the Index Score, goal dashboard, and trend dashboard for all indicators and goals, spillover scores, raw values, normalized scores, dashboard ratings, trends, and goal scores. The SDG Index Score and all other indicators are retroactively calculated across time using time series data. Missing data are treated by carrying forward

time series data. Data come from international and rigorous organizations that operate extensive validation processes, like the World Bank, the Organisation for Economic Co-operation and Development, the World Health Organization, the Food and Agriculture Organization, and the International Labour Organization. Therefore, the dataset characteristics limit any biases related to single-source utilizations, and the complete coverage of UN members avoids sample selection issues.

2.2.2 The Bass Diffusion Model

The Bass diffusion model (Bass, 1969) helps in understanding innovation adoption processes within populations by analyzing commercial product uptake through its characteristic S-shaped adoption curve. This temporal pattern reflects the general transition from slow adoption to rapid acceleration and eventual market saturation across diverse innovation contexts. The model's theoretical foundation rests on two complementary adoption mechanisms: innovation-driven adoption (external influence parameter) and imitation-driven adoption (internal influence parameter). It is important to situate the Bass model (BM for short) in relation to innovation diffusion theory (Rogers, 2003), which conceptualizes the adoption of innovations through five adopter categories (innovators, early adopters, early majority, late majority, laggards) and emphasizes the role of communication channels, social systems, and time in shaping diffusion patterns. Rogers' framework is highly influential in providing a qualitative, sociological account of how innovations spread; however, it does not provide a direct quantitative mechanism for forecasting adoption dynamics or for estimating parameters from empirical data. By contrast, the BM translates Rogers' qualitative insights into a parsimonious mathematical formulation. This quantitative specification allows researchers to fit empirical time series, estimate country-specific parameters, and make predictions about saturation dynamics, which is essential for the analysis of SDG trajectories. The adaptation of the Bass diffusion model to SDGs represents a paradigmatic shift toward complex socioinstitutional transformation processes (Wonglimpiyarat, 2025). As such, this recontextualization requires careful consideration of the unique characteristics that distinguish SDG implementation from traditional innovation diffusion. The innovation parameter in the SDG context encompasses exogenous drivers, including governmental policy frameworks, complex international agreements and multilateral initiatives, and institutional mandates that independently promote SDG-aligned practices. These

external influences are manifested through mechanisms such as renewable energy subsidies, carbon pricing policies, environmental regulations, and international development assistance programs. The imitation parameter captures endogenous social dynamics, including policy coherence, knowledge spillovers, normative convergence, and complex cross-border learning processes, that increase SDG implementation through peer effects. This parameter encompasses regional adoption of successful sustainability models, South–South cooperation mechanisms, and knowledge transfer networks that facilitate the diffusion of best practices (Collste et al., 2017). The market potential translates to SDG saturation potential, i.e., the maximum feasible level of sustainable development achievement a country can attain given its structural constraints (Marzouk et al., 2022). These constraints may include economic capacity, institutional quality, natural resource endowments, demographic characteristics, and geopolitical complexity and stability. Hence, the adapted Bass diffusion model provides a theoretically grounded foundation for modeling SDG achievement trajectories across heterogeneous country contexts (Fu et al., 2020). Its reformulation – incorporating dynamic saturation potentials, country-specific heterogeneities, and complex system feedback – enables robust predictions of SDG saturation timing and identifies leverage points for accelerating global sustainability transitions.

Mathematically, the Bass model reads:

$$\frac{dF(t)}{dt} = (p + qF(t))(M - F(t)), \quad (2.1)$$

where originally $F(t)$ denoted the cumulative fraction of adopters at time t ; p represented the proportion of adopters driven by forces exogenous to the social dynamics of diffusion, such as advertisement or personal inclination towards innovation; and q represented the proportion of adopters driven by social pressure, i.e., the fact that many other people have adopted. The parameters p and q are referred to as an innovation parameter and an imitation parameter, respectively. Parameter M is the market potential, i.e. the total amount of adoption registered at the end of the process. The solution to this differential equation describes the S-curve, characterized by an initial period of slow adoption dominated by innovators, followed by rapid adoption driven by imitators, and finally a saturation phase when M is approached. The Bass model has traditionally been applied in contexts such as consumer product adoption, technology diffusion, and innovation management. However, the overall flexibility of innovation diffusion models and theories, enables adaptation to other domains

where a diffusion-like process occurs (Wonglimpiyarat, 2025). In this chapter, we have applied the Bass framework to data related to the SDGs for a large number of countries. The S-curve generated by the BM seems well suited to deriving insights into this process. Similarly to innovation diffusion, national efforts toward SDG targets often begin slowly due to limited resources, institutional inertia, or other structural barriers, but over time, as progress accelerates through increased investment, policy implementation, and knowledge sharing, the rate of improvement may rise sharply before eventually stabilizing as the full potential of the country is approached.

2.2.3 The Adaptation of the Bass Diffusion Model

To carry out this analysis, we fitted the Bass curve (equation 2.1) to the time series of the SDG results for each country. The dataset includes data on the aggregate SDG Index Score as well as indicators for the 17 different goals. The database also contains highly granular data with time series for single measures contained in each goal, but no analysis has been carried out for these subgoals. The time span of the dataset covers 22 years, from 2000 to 2021, with one data point for each year corresponding to the level of fulfillment of the SDG on a scale ranging from 0 to 100. With respect to the SDG Index Score time series, we performed a curve fitting of the S-curve produced by the BM using nonlinear least squares. In particular, this means searching for the set of parameters (p, q, M) that produces the S-curve that minimizes the squared difference between the curve itself and the data points of the time series. It should be noted here that the canonical Bass has null initial condition and grows to M . Hence, to apply the Bass framework in this case, we had to shift the SDG Index Score properly and later readjust the fulfillment potential level M . Depending on the case and on the quality of the fitting, some precautions have been adopted to ensure the best fitting of the data points. First, a Monte Carlo exploration of the space of the parameters was performed in case of an unsatisfactory result of the fitting for stiffer cases. In some cases, the high fluctuations in the data resulted in poor fitting of the curve. In these cases, a moving average smoothing or an exponential smoothing was performed prior to the application of the least squares algorithm, depending on which led to the best performance. For a minority of countries, the best fitting curve seems to underestimate the actual curve, likely due to idiosyncrasies in the trajectories of the Index Score. For these countries,

we applied a case-specific fine-tuning of the parameters—a common practice when fitting Bass curves to real data. For a small subset of countries, the time series appears highly linear. In these cases, the least squares method over the full set of parameters (p , q , M) led to a meaningless overestimation of parameter M , which far exceeded the cap value of 100. For these countries, we forced the market potential to 100. Moreover, for some countries, the trajectory of the SDG Index Score does not follow an S-shaped growth. For these nations, it is pointless to apply the described method as one cannot retrieve useful insights from the Bass framework. For each country, after having applied the corrections and assessed the best fit parameters, the time of the peak in the performance and the expected saturation time were computed. The former refers to the time—which can be either in the past or in the future—when the maximum increase in the year-to-year score is detected. Since this refers to the Bass S-curve rather than the real data, it is not influenced by random fluctuations in the dataset. The latter refers to the time in which the country is expected to reach a satisfactory proportion of its fulfillment potential. In this work, the satisfaction threshold for computing the saturation time is 100%, but of course, this can be modified to be less demanding. This method can be applied in the same way to the time series for the single goals for the countries in the dataset. Due to these case-by-case exceptions, though, doing so is time-consuming. As a preliminary study, we report in the next section some results derived from the application of this method only to the aggregated SDG Index Score. If needed, more granular analysis can be performed.

2.3 Results

Using the SDG Index Scores as the target series, the Bass curve provides a compact and informative description of national progress for the bulk of entities. Evaluated by the RMSE normalized by each country's observed range, fits are typically tight: many countries land well below 0.10, with a large mass clustered between roughly 0.03 and 0.08. In practical terms, average errors are on the order of only a few percent of each country's total movement over 2000-2021. That is consistent with what the eye sees in the plots, where the fitted S-curve tracks the slow-start/acceleration/saturation pattern closely. A smaller group sits in the 0.10–0.15 band, where phase/timing errors become visible but the macro shape is still captured. Only a minority (30

out of 177 entities) exceed 0.15, and a very small set (13) crosses 0.20; these latter cases are the ones where the S-curve is visibly deviating from the data, or where the series exhibits shocks and stalls that a single-regime curve is not designed to follow. Four countries (Azerbaijan, Cuba, Gabon and Luxembourg) exhibit invalid Bass parameters and a concave improvement. Figure 2.1 shows the plot of the Bass curve fitted to the SDG Index Score for four notable cases. In particular we show the plots of China as the best fitting case (2.1a), USA as an example of fit with moving average smoothing (2.1b), Japan as an example of fit with fulfillment potential M fixed (2.1c), and Venezuela as an example of country that does not follow an S-curve in its SDG Index Score progression (2.1d).

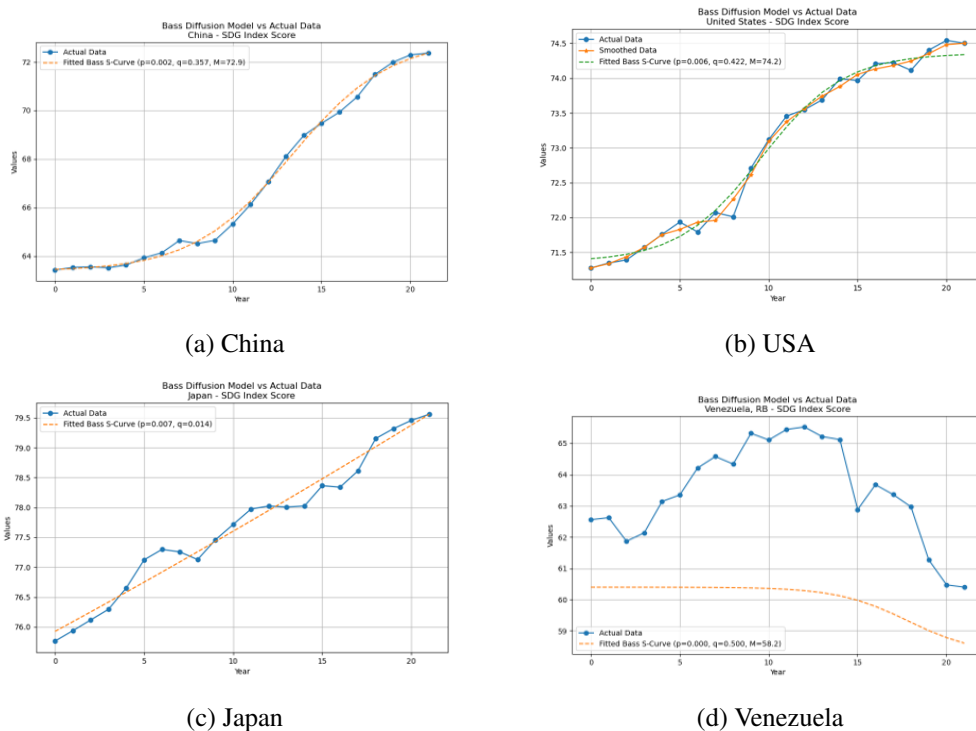


Fig. 2.1 SDG Index Score progression and fitted Bass curve for four countries between years 2000 and 2021.

Aggregates behave as expected. Regional and income-group composites – for example, Middle East & North Africa, OECD members, high-income countries – display very low normalized errors, often around two to three percent of their observed range. This is the signature of idiosyncrasy smoothing: aggregation mutes local year-to-year noise and irregularities, revealing a clean S-shaped backbone that the Bass curve can summarize with little ambiguity. Figure 2.2 shows the plots

for these aggregates. By contrast, individual countries with volatile trajectories or structural breaks (conflict, policy shocks, measurement changes) are over-represented among the weakest fits; the very high RMSE values for some fragile states and crisis economies reflect genuine departures from a single smooth diffusion path rather than a mere estimation defect. Table 2.1 contains the estimated parameters and errors for the plots shown in figures 2.1 and 2.2. The whole set of results is summarized in appendix A.

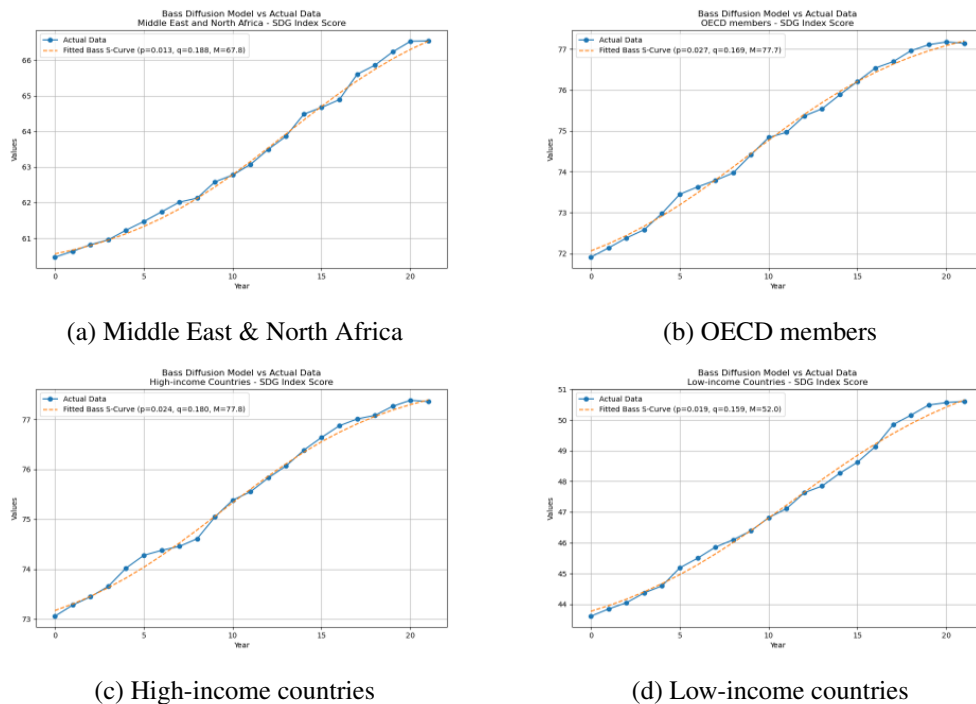


Fig. 2.2 SDG Index Score progression and fitted Bass curve for four aggregates between years 2000 and 2021.

2.4 Discussion

Clustering by socioeconomic status and by geography yields complementary angles on SDG diffusion. Socioeconomic tiers – high-, middle-, and low-income – sort countries by the very capacities that condition diffusion: economic capacity, institutional quality and complexity, and human capital. In those tiers we see predictable differences in pace, consistent with work showing that innovation ecosystems and SDG progress are mediated by wealth, education, and institutional robustness, which

Entity	p	q	M	Peak	Saturation	RMSE
China	0.0025	0.3571	72.9148	2013	2033	0.0222
Japan	0.0067	0.0144	100.0000	2036	2121	0.0493
United States	0.0056	0.4218	74.2381	2010	2024	0.0550
Venezuela	0.0001	0.5000	58.2483	2018	2000	0.7109
High-income	0.0239	0.1799	77.7987	2009	2032	0.0229
Low-income	0.0192	0.1589	52.0185	2011	2039	0.0246
M.E. & N.A.	0.0125	0.1879	67.8083	2013	2039	0.0204
OECD	0.0266	0.1693	77.7033	2009	2032	0.0215

Table 2.1 Estimated Bass parameters, relative peak and saturation years and RMSE for the countries and aggregates in figures 2.1 and 2.2

in turn shape both “innovation” and “imitation” effects (Nylund et al., 2021). The heterogeneity in Bass parameters across these groups is therefore not a curiosity but an expected reflection of foundational constraints and advantages. Geographic clustering, by contrast, picks up cohesion and spillovers: shared cultural, political, and environmental contexts that ease knowledge transfer and normative convergence (Owen and Vedanthachari, 2022). Yet large within-region spreads persist, underlining that proximity does not dissolve structural barriers such as resource gaps, fragmented and complex governance, and inequality still dominate outcomes (Grama-Vigouroux et al., 2023). This echoes the literature’s warning that ecosystem collaboration often runs into implementation frictions where capacities are uneven and institutions are weak (Berry et al., 2024; Chaparro-Banegas et al., 2024). Regional cooperation (SDG 17) can accelerate diffusion through partnerships and policy harmonization (Owen and Vedanthachari, 2022), but the binding constraint remains socioeconomic. These patterns align with our fit diagnostics. High-income countries and OECD members, coincident as a socioeconomic cluster, show stronger Bass fits and higher apparent saturation potential, consistent with advanced innovation ecosystems and dense policy infrastructures (Husainy et al., 2024; Jütting, 2022). Many low-income countries remain in slower-diffusion clusters even when adjacent to better performers, in line with documented digital divides and resource constraints that widen SDG disparities (Chaparro-Banegas et al., 2024; Yuan et al., 2023). The upshot is not that ecosystems are optional; rather, their effectiveness on climate/energy (SDGs 7, 13), circularity (SDG 12), and inclusion (SDG 10) is conditional on overcoming those structural barriers (Brás and Robaina, 2024; Song, 2025). Clustering patterns also register the spatial unevenness of innovation benefits. Urban-centric hubs concen-

trate capital, skills, and complementary assets, while rural and marginalized areas are left with thinner infrastructures and weaker connectivity. These mechanisms reproduce spatial inequality even as national averages improve (Yuan et al., 2023). In low-income clusters, digital ecosystems can unintentionally deepen exclusion: infrastructural deficits restrict access and slow the imitation channel that would otherwise diffuse sustainable practices beyond early adopters (Chaparro-Banegas et al., 2024). Misalignment between national innovation priorities and local SDG needs compounds the problem: policy that privileges generic competitiveness over context-critical goals such as water security or biodiversity can yield measurable gaps in places where resources are scarce and trade-offs bite hardest (Albitar et al., 2023).

These results echo the dual nature of innovation ecosystems emphasized in recent work. On the one hand, ecosystems enable systemic collaboration (SDG 17), the scaling of green technologies (SDG 9), and adaptive resilience (SDG 3) by coordinating multi-stakeholder efforts, mobilizing digital tools, and aligning incentives. On the other hand, they can amplify disparities (SDG 10) and generate environmental or social trade-offs (SDGs 7, 15), particularly where capabilities and safety nets are thin. A salient illustration is the negative correlation reported between some green innovation waves in G7 economies and poverty alleviation, as high capital costs crowd out social investment – an aggregate pattern consistent with the heterogeneity we observe in diffusion clusters (Islam, 2025). In sum, the clustering evidence points to an interaction rather than a hierarchy: socioeconomic capacity sets the feasible envelope for diffusion; geography and regional cooperation shape spillovers within that envelope; and governance choices determine whether ecosystems translate capability into equitable SDG progress (United Nations, 2015). Effective policy therefore has to move on both margins: reinforce foundational capacities (infrastructure, inclusion, local problem-fit) while orchestrating regional mechanisms that lower coordination costs and propagate workable solutions across borders.

2.5 Conclusions

This chapter demonstrates that the Bass diffusion model can be effectively adapted to analyze national progress toward Sustainable Development Goals, capturing

the characteristic S-shaped adoption curves of SDG achievement across diverse countries. The model's parameters reveal significant heterogeneity influenced by socioeconomic and geographic factors, reflecting the complex dynamics of innovation ecosystems. While high-income countries tend to approach saturation earlier, lower-income nations face structural constraints limiting their SDG fulfillment potential. These findings highlight the importance of innovation ecosystems and governance in accelerating sustainable development. Overall, the Bass model provides a valuable quantitative framework for forecasting SDG saturation timing and informing policy interventions.

The application of the Bass diffusion model in this study extends its traditional use from product and technology adoption to the complex domain of achieving SDGs at the national level. The model's foundational premise—that adoption dynamics arise from two complementary mechanisms, namely innovation-driven external influences and imitation-driven internal social interactions—offers a parsimonious yet powerful lens through which to understand the nonlinear progression of SDG fulfillment. By reinterpreting the market potential parameter as the SDG saturation potential, this adaptation accounts for structural constraints and capacities that are unique to each country, such as economic resources, institutional quality, and geopolitical stability. Theoretically, this reformulation bridges innovation diffusion theory with complexity science by acknowledging that SDG progress emerges from the interplay of multiple heterogeneous agents within national innovation ecosystems. These ecosystems are characterized by adaptive, nonlinear feedback loops among governments, firms, and civil society actors, which influence the acceleration and eventual saturation of sustainable development efforts. The Bass model's S-shaped curve effectively captures the transition from initial slow progress – often hindered by institutional inertia and resource limitations – to rapid acceleration driven by policy coherence, knowledge spillovers, and normative convergence, culminating in a plateau as countries near their fulfillment ceilings.

Moreover, the model's parameters provide insight into the relative strength of external policy drivers versus internal social dynamics across different socioeconomic and geographic clusters, highlighting the importance of context-dependent governance frameworks. This aligns with the emphasis of complexity theory on emergent behaviors arising from microlevel interactions within complex adaptive systems. While the BM simplifies some aspects of these interactions by aggregating adopters, its successful application here suggests that even relatively simple diffusion

models can yield meaningful insights into complex socioinstitutional transformations when carefully adapted. Theoretically, this study invites further integration of the Bass framework with network-based and agent-based models that explicitly capture heterogeneity and multilevel interactions, enhancing the understanding of how innovation ecosystems evolve and influence SDG trajectories. It also underscores the need to consider dynamic saturation potentials and threshold effects to better reflect the evolving capacities and constraints within innovation ecosystems. In sum, this work contributes to the theoretical literature by demonstrating the Bass model's flexibility and relevance for modeling complex, multi-actor sustainability transitions in national contexts.

From a practical standpoint, the adapted Bass diffusion model offers policy-makers, innovation managers, and development practitioners a quantitative tool for forecasting the timing and extent of national SDG achievement, enabling more informed strategic planning. By estimating the innovation and imitation parameters, decision-makers can discern the relative impact of external policy initiatives versus internal social dynamics, guiding the allocation of resources and the design of interventions to accelerate sustainable development. For instance, a higher innovation parameter suggests that strengthening external drivers – such as regulatory frameworks, subsidies for green technologies, and international cooperation – can effectively stimulate early SDG adoption. Conversely, a higher imitation parameter highlights the critical role of endogenous factors like peer learning, knowledge spillovers, and regional cooperation networks, suggesting that fostering collaborative governance and multi-stakeholder partnerships can increase progress. Understanding these dynamics allows governments to tailor policies to their specific contexts, focusing on either enhancing policy incentives or facilitating social diffusion mechanisms. The ability of the model to predict saturation times also assists in setting realistic targets and monitoring progress, as it helps to identify countries or regions at risk of stagnation due to structural barriers or governance challenges. This insight supports the prioritization of capacity-building efforts, institutional reforms, and investment in innovation ecosystems where they are most needed. Furthermore, recognizing the heterogeneity of diffusion patterns across socioeconomic clusters enables international organizations and donors to customize support strategies, addressing disparities between high-income and low-income countries. In addition, the Bass model's framework can inform private sector actors and innovation hubs by signaling market readiness and potential demand for sustainable technologies,

thereby optimizing investment decisions and scaling strategies. It also aids in anticipating the impact of disruptive events or policy shocks on SDG trajectories, enabling the use of adaptive management in complex and uncertain environments. Overall, this practical application of the Bass diffusion model equips stakeholders with a robust, data-driven approach to enhance coordination, improve policy coherence, and accelerate the diffusion of sustainability innovations within national innovation ecosystems, ultimately advancing the global 2030 Agenda.

2.5.1 Limitations and Future Research Directions

Despite its valuable insights, this study acknowledges several limitations inherent in applying the Bass diffusion model to complex socioinstitutional phenomena like SDG achievement. The model's assumption of homogeneous adopter populations and constant parameters over time simplifies the diverse and evolving nature of national innovation ecosystems. Some countries exhibited poor model fit or anomalous parameter estimates, reflecting sociopolitical instability or data quality issues that the model cannot fully capture.

Future research should explore methodological enhancements by integrating the Bass framework with agent-based and network diffusion models to explicitly represent heterogeneous actors, multilevel governance structures, and dynamic feedback loops. Incorporating time-varying parameters would enable the model to better reflect changing policy environments, shocks, and adaptive behaviors. Additionally, leveraging richer, higher-frequency data sources, such as real-time policy implementation metrics or social network analyses, could improve parameter estimation and model validation. Further empirical studies could extend the analysis to subnational levels or specific SDGs to uncover more granular diffusion patterns and contextual factors. Investigating the interplay between innovation ecosystems and negative externalities, such as inequality or environmental trade-offs, would also deepen understanding of the “dark side” of innovation. Future research efforts should also be devoted to coordination-level efforts in platform ecosystems and innovation ecosystems in general, with a view to grand challenge resolution. Finally, interdisciplinary approaches combining complexity science, behavioral economics, and sustainability transitions theory hold promise for refining theoretical frameworks and developing more comprehensive models to guide policy and practice in achieving sustainable development.

Chapter 3

The statistical mechanics of innovation diffusion

Chapter 2 looked at the observable outcome of a complex system – the SDG Index as produced by national innovation ecosystems – through a wide-angle lens. Using a compact S-curve model, it summarized country trajectories and offered comparable diagnostics across contexts. That choice was deliberate – such curves travel well – but it left the underlying machinery implicit. Here, we return to the original landscape of the diffusion of an innovation and we open the black box. We start from simple behavioral rules for many interacting, boundedly rational agents and show how their local updates aggregate into a tractable adoption equation. In doing so, the levers that matter for diffusion remain explicit rather than folded into a pair of parameters, and the familiar S-shape becomes an emergent outcome rather than a premise. The discussion in this chapter follows the paper from Masali et al. (2025).

This shift from description to mechanism is also a shift in what we can learn. A micro-founded law allows disciplined estimation and uncertainty assessment, and it clarifies which forces – information, experience, heterogeneity – are needed to reproduce observed paths. It also sets up the next step in the thesis: chapter 4 will build on the model introduced here adding expectations and technological performance to explain rise–peak–trough dynamics.

3.1 General Context

In the operations research literature, to which this chapter primarily refers, the Bass model (Bass, 1969) remains the most widely adopted framework for modeling new product diffusion. It counts hundreds of successful applications ranging from satellite TV to smartphones (Bass, 2004), from medical technologies (Simon and Sebastian, 1987) to commerce platforms (Li et al., 2020). As explained in section 2.2.2, the Bass model divides buyers into two categories: *innovators* and *imitators*. Innovators purchase the product independently of others, typically driven by intrinsic interest or exposure to advertising. These people are pivotal in the early stages of the diffusion process, since their purchase does not depend on previous sales of the product. Imitators buy the product under social pressure—whether framed as word-of-mouth, contagion, or internal influence. This pressure grows with the number of buyers, becoming more and more dominant as time goes by. Being the population limited, the process then slows down when approaching the saturation point.

The BM is an aggregate model: it starts from the observed macro-pattern (typically an S-shaped adoption curve) and models the phenomenon accordingly. Aggregate models are favored by managers and decision-makers for their simplicity and interpretability. They provide a compact, analytical summary of sales dynamics and can be estimated from readily available aggregate data. However, they rely on strong assumptions on the underlying behavioral processes, and they are less equipped to capture heterogeneity or psychological nuance at the individual level. Despite the descriptive power of the BM and its many variants, it remains an aggregate-level approach, often disconnected from individual-level behaviors. Our work in this chapter aims to fill this gap by deriving an aggregate adoption curve from micro-level rules, grounded in agent interaction dynamics.

Agent-based models are based on a bottom-up approach that can overcome some of the aforementioned shortcomings. In AB modeling, no assumptions are made on the whole system: individual-level rules are set for the agents and the population-level behavior emerges from the interactions of those agents. This makes it possible to model differences in behavior, preferences, and psychological traits. However, most ABMs in the literature are simulation-based, which limits their analytical tractability. They typically simulate a population of agents over time and track adoption at each step, which can require significant computational resources. Moreover, calibration can be difficult because the available data is often aggregate,

while the model parameters operate at the micro-level. In this chapter, we propose an agent-based new-product diffusion model that combines the flexibility of agent-based models with the analytical clarity of aggregate models. Our approach is based on a version of the Boltzmann equation, originally developed in kinetic gas theory, but adapted to describe social interactions. The idea is to draw a parallelism between the molecules interacting in a gas and the people interacting in society, in order to be able to use the same mathematical formalism. When molecules collide, they exchange kinetic energy following the laws of mechanics and this modifies their position and velocity, i.e. their *microscopic state*. In this model, we set the microscopic state of our agents to be the perceived value of the new product. When people interact with each other (or with the advertisement campaign), they exchange information and change their perceived value. We will assume people follow a *compromise dynamics*: when they interact, they move partially toward each other's beliefs.

This modeling approach offers several advantages:

- It adopts a bottom-up perspective, linking aggregate outcomes to micro-level behavioral rules.
- It performs competitively against state-of-the-art extensions of the Bass model (Cosguner and Seetharaman, 2022).
- Its parameters are directly interpretable in terms of marketing levers (e.g., price, advertising intensity), making it useful for managerial decision-making.
- It provides a flexible framework: by adjusting the agent-level rules, one can model repeat purchases, abandonment, segmentation, and more—without needing to assume a specific functional form for the aggregate adoption curve.

The chapter is organized as follows. In the next section, we review the literature on Bass-type aggregate models, agent-based models, and the use of kinetic theory in social systems. Section 3.3 describes the Boltzmann equation and the kinetic framework that are the foundations of the model. Later, section 3.4 presents the kinetic diffusion model and section 3.5 evaluates its performance. Section 3.6 concludes the chapter.

3.2 Related Literature

This section provides a focused review of the innovation diffusion literature in operations research. For a broader overview in sociology and economics, see Rogers (2003). The success and wide range of applications of the BM are clear if we look at the number of extensions developed to address a range of challenges marketers face when launching new products. These include, but are not limited to, product substitution, pricing, and advertisement calibration. For a complete review of the literature on new-product diffusion models up to 2010 see Peres et al. (2010).

The first adaptation of the BM to product substitution appeared in Norton and Bass (1987) and was later refined in several improved versions (Jiang and Jain, 2012; Singh et al., 2012). The problem of dynamic pricing for new products has been studied both within and outside the BM framework – see, for instance, Bass and Bultez (1982), Kalish and Lilien (1983), Dockner and Jørgensen (1988). Because the original BM lacks explicit marketing levers, it was extended by Bass et al. (1994) into the Generalized Bass Model (GBM), which incorporates price and advertising as part of the marketing mix. Since then, many versions of the GBM have appeared to address the various decision-making needs, particularly around pricing and advertising. Important contributions in this area include Robinson and Lakhani (1975), Krishnan et al. (1999) and Krishnan and Jain (2006). Other aspects of the diffusion process that have been studied inside the BM framework are its dependence on opinion distribution in a population (Fan et al., 2017; Han et al., 2022) and supply restrictions (Ho et al., 2002; Kumar and Swaminathan, 2003; Shen et al., 2011). For the sake of completeness let us stress that both these topics have been extensively studied outside the BM framework as well, see as examples Assenova (2018) and Iyengar et al. (2011) for opinion-based models and Keith et al. (2017) for supply constraints. Some authors have criticized the GBM as a prescriptive model for marketing strategy (Bass et al., 2000; Fruchter and Van den Bulte, 2011). The recent work by Cosguner and Seetharaman (2022) overcomes some of these limitations using a utility-based approach.

Because the BM is an aggregate model, it requires strong assumptions about system behavior. This has motivated efforts to disaggregate the diffusion process, by modeling individual-level adoption and deriving the resulting aggregate patterns. These efforts went through different paths: Sinha and Chandrashekar (1992) used hazard models to incorporate different timing in adoption for different buyers; a

spatial dimension is taken into account in the works by Bronnenberg and Mela (2004) and Bell and Song (2007) while the work by Li et al. (2014) is based on network theory. Our approach falls under the broader class of agent-based models, which trace back to Goldenberg et al. (2000). Aggregate models and ABMs represent two opposite approaches to the same problem and can be used in parallel to satisfy different needs (Rahmandad and Sterman, 2008). The vast majority of ABMs for new-product diffusion are simulation-based, built to explore the diffusion of specific products or behaviors in particular market contexts. Notable applications include Palmer et al. (2015) on private photovoltaic in Italy; Stummer et al. (2015) on repeat purchase modeling; Bastani et al. (2016) on energy-savings behaviors, Zhang et al. (2022) on photovoltaic adoption in Singapore; Rotaris and Scorrano (2023) on car-sharing diffusion in an Italian region. Important exceptions to simulations in agent-based modeling exist. Fibich and Gibori (2010), for example, aggregate agent dynamics on a network to derive analytical expressions for product diffusion. Another promising direction is the mean-field game (MFG) approach, which leverages game theory to model strategic multi-agent behavior. In this stream of literature, Chenavaz et al. (2021) study price optimization in a multi-firm new-product diffusion process; while Chaab et al. (2022) imagine a Stackelberg game between a firm and a large number of strategic consumers and draw the diffusion dynamics from the mean-field approximated Nash equilibrium.

In this chapter, we adopt an alternative kinetic approach grounded in the mathematical formalism of statistical mechanics. Inspired by the kinetic theory of gases, this approach models the population as a system of interacting agents whose beliefs and behaviors evolve over time—analogue to molecules exchanging energy upon collision. By treating interactions as belief-updating events, we derive closed-form differential equations for aggregate quantities from individual-level rules. For a comprehensive introduction to this methodology, see Pareschi and Toscani (2013). In recent years, the kinetic framework has been successfully employed to describe a wide variety of social and biological phenomena. Of particular interest for this case are the applications regarding opinion formation (Düring and Wright, 2022; Toscani, 2006) and epidemiology (Della Marca et al., 2021; Loy and Tosin, 2021).

3.3 Kinetic Framework

In this section, we reinterpret the classical mathematical framework of the Boltzmann equation from the perspective of operations research and economics. Our goal is to model how individual beliefs about the value of a new product evolve due to interaction and information flow. Inspired by kinetic theory, which studies how microscopic interactions in gases lead to macroscopic behaviors, we borrow the structure of those models and apply it to agent-based settings where individuals update their beliefs through social and media interactions.

Rather than modeling particles in a gas, we view our agents as consumers or market participants who revise their evaluation of a product over time. This revision is influenced by peer interactions (e.g., word-of-mouth), background signals (e.g., advertising), and stochastic fluctuations that represent individual differences or unmodeled factors.

3.3.1 The Boltzmann equation

Let $f(t, v) \geq 0$ denote the distribution of agents over perceived product values v at time t . The function f can be interpreted as a probability density over a continuum of belief states in the market, in the sense that the quantity $f(t, v)$ indicates the number of agents with perceived value $v \in [v, v + dv]$ at time t . Its evolution is modeled through an interaction-based dynamic equation inspired by models of collective behavior, namely the Boltzmann equation (Boltzmann, 1872), which, in this context, captures the effect of interactions on agents' beliefs:

$$\frac{\partial f(t, v)}{\partial t} = \mu Q(f, f)(t, v). \quad (3.1)$$

Where $Q(f, f)(t, v)$ represents the effect of pairwise interactions on the belief distribution, and μ is the frequency of the interactions. The evolution of f is driven by two components:

- Peer influence: when two agents interact, they influence each other's perception of the product.
- External signals: agents may also adjust their perception based on external campaigns or real-world experience.

An important feature of the interaction mechanism is that it preserves the total number of agents. This implies:

$$\int_V Q(f, f)(t, v) dv = 0 \quad \forall t. \quad (3.2)$$

This property allows us to give the density $f(t, v)$ a probabilistic sense. If we start with a density $f(0, v)$ such that $\int_V f(0, v) dv = 1$, we will be sure that this will continue to be true for all subsequent times. Hence

$$\int_V f(t, v) dv = 1 \quad \forall t > 0. \quad (3.3)$$

Using this, we can compute expectations of relevant market quantities. Let $\varphi(v)$ denote any measurable characteristic of an agent as a function of their valuation. Then the average value of $\varphi(v)$ at time t is given by:

$$\langle \varphi(v) \rangle = \int_{\mathbb{R}} \varphi(v) f(t, v) dv, \quad (3.4)$$

where we omit the dependency on time t to lighten the notation. Two specific statistics will be central in our analysis:

- $\rho(t) = \int f(t, v) dv$: the total number of agents in the population,
- $m(t) = \int v \cdot f(t, v) dv$: the average perceived value of the product across the population.

Note that it is not trivial to compute the number of agents because we will employ a label switch mechanism: while the whole population remains constant, the population associated with the different labels will change. For a comprehensive discussion on the Boltzmann equation see Cercignani (1988).

3.3.2 Binary interactions

The interactions correspond to every moment in which value v may change due to external inputs, namely word-of-mouth, advertising or actual usage. We assume that when two agents interact, their perceived values shift toward a compromise. This mechanism reflects bounded rationality: agents cannot compute optimal decisions in

isolation, so they rely on peer comparisons. Mathematically, if v and v_* represent the pre-interaction valuations, the post-interaction value is:

$$\begin{aligned} v^{post} &= I(v, v_*) \\ v_*^{post} &= I(v_*, v), \end{aligned} \quad (3.5)$$

where $I(\cdot, \cdot)$ is the interaction function encoding how beliefs are exchanged or adjusted during interaction. It will be symmetric and stochastic. To capture the dynamics of this belief-updating process at the population level, we consider the effect of a small time interval Δt , during which interactions may or may not occur. We assume that a binary interaction takes place with probability $\mu\Delta t$, where μ represents the interaction rate between the agents. This is modeled using a Bernoulli random variable, independent of the specific content of the interaction. Under this setup, an agent's valuation is updated with probability $\mu\Delta t$ and left unchanged otherwise:

$$v' = I(v, v_*) \cdot \mu\Delta t + v \cdot (1 - \mu\Delta t), \quad (3.6)$$

and hence

$$\langle \varphi(v') \rangle = \langle \varphi(I(v, v_*)) \rangle \cdot \mu\Delta t + \langle \varphi(v) \rangle (1 - \mu\Delta t). \quad (3.7)$$

Re-organizing the members and computing the limit for small time intervals it is easy to prove the following equality (Pareschi and Toscani, 2013):

$$\frac{d}{dt} \langle \varphi(v) \rangle = \mu (\langle \varphi(v') \rangle - \langle \varphi(v) \rangle). \quad (3.8)$$

If we explicit the average through equation (3.4) we obtain

$$\frac{d}{dt} \int_V \varphi(v) f(t, v) dv = \mu \left\langle \iint_{V \times V_*} \left(\varphi(v') - \varphi(v) \right) f(t, v) f(t, v_*) dv dv_* \right\rangle, \quad (3.9)$$

where we have assumed statistical independence between the interacting agents—i.e., the joint distribution $f^{(2)}(t, v, v_*)$ factorizes as $f(t, v)f(t, v_*)$.

To reflect heterogeneity in information processing and judgment, we introduce a stochastic noise term into the interaction function $I(\cdot, \cdot)$. This captures private signals, unobservable traits, and random deviations from strict imitation or compromise. With a change in the notation, $\langle \cdot \rangle$ indicates now the expected value over the stochastic interaction function.

This represents the weak form of the Boltzmann equation. The meaning of this

formula is that the variation in time of the average of an observable quantity in an interacting system is equal to the mean variation of that quantity in the interactions.

Interaction with a background

Agents are not influenced only by peers, they are also exposed to marketing signals. We model this using a background distribution $G(v)$, representing the message conveyed by an advertising campaign. The strength and content of the campaign are encoded by the weight and position of peaks in $G(v)$.

$$G(v) = \sum_{i=0}^n w_i \delta(v - v_i), \quad (3.10)$$

where w_i (such that $\sum_{i=0}^n w_i = 1$) are the weights of information and $\delta(\cdot)$ is the Dirac delta function (Hassani and Hassani, 2009). This background influence can be incorporated in the same interaction formalism, treating it as a one-sided interaction with a fixed signal:

$$\frac{d}{dt} \int_V \varphi(v) f(t, v) dv = \mu \langle \iint_{V \times V_*} (\varphi(v') - \varphi(v)) f(t, v) G(v_*) dv dv_* \rangle. \quad (3.11)$$

Although $G(v)$ is assumed to be fixed, it can be generalized to a time-dependent function to reflect dynamic advertising strategies.

3.3.3 Label switch process

To model the decision to adopt or not adopt the innovation, we divide the population into two groups: *users* and *non-users*. Agents transition from non-user to user status when their perceived value exceeds the product price. This is formalized through a Markovian label-switching process similar to Loy and Tosin (2020), where the transition probability depends on individual valuation. We equip each agent with a label $i \in \mathcal{I} = 1, 2$ and define a transition event as an event where a purchase may occur. This will happen according to frequency λ . If such an event occurs, the transition probability from one label to the other will be:

$$T(v; i|j) \in [0, 1] \quad \forall v \in V, i, j \in \mathcal{I}. \quad (3.12)$$

We have defined a Markov process where the transition probability is belief-dependent. To ensure the process is well-defined and the total population remains constant, we require:

$$\sum_i T(v; i|j) = 1 \quad \forall v \in V, y \in \mathcal{J}. \quad (3.13)$$

If we plug this into the kinetic framework for a distribution $f(t, i, v)$ of agents with label i , we obtain:

$$\sum_{i=1}^n \varphi(i) \frac{d}{dt} f_i(t) = \lambda \sum_{i=1}^n \sum_{j=1}^n (\varphi(i) - \varphi(j)) T(i|j) f_j(t), \quad (3.14)$$

where the subscript indicates the distribution restricted to some label. If we choose $\varphi(i) = 1$ for $i \in \mathcal{I}$ and $\varphi(j) = 0 \quad \forall j \in \mathcal{J} \setminus i$ we are counting agents that belong to group i . Thus, we can compute the evolution of density for group i as

$$\frac{d}{dt} f_i(t) = \lambda \left(\sum_{j=1}^n T(i|j) f_j(t) - f_i(t) \right) \quad i = 1, \dots, n. \quad (3.15)$$

We now combine this label-switching mechanism with the belief-updating interactions described earlier, leading to a joint evolution of agents' beliefs and their adoption status. Let us take distribution function $f_i(t, v)$, we will have, from mass conservation:

$$\sum_{i=1}^n \int_V f_i(t, v) dv = 1 \quad \forall t > 0. \quad (3.16)$$

The addends of the sum are the masses of the agents with label i :

$$\rho_i(t) := \int_V f_i(t, v) dv. \quad (3.17)$$

If we consider the label switch process and the interaction process (both with other agents and with the background) happening in parallel with frequencies λ and μ , respectively, and with a time interval Δt small enough to neglect the second order contributions we will obtain a non-conservative weak Boltzmann equation for a

single group i of the form:

$$\begin{aligned} \frac{d}{dt} \int_V \varphi(v) f_i(t, v) dv &= \lambda \int_V \varphi(v) \left(\sum_{j=1}^n T(v; i|j) f_j(t, v) - f_i(t, v) \right) dv \\ &+ \mu \sum_{j=1}^n \iiint_{V^3} \langle \varphi(v') - \varphi(v) \rangle f_i(t, v) f_j(t, v_*) G(\hat{v}) dv dv_* d\hat{v} \quad (3.18) \end{aligned}$$

In this case, the interaction rule is the same for every agent, regardless the group it belongs to. From here, it is easy to describe a situation where the interaction rule depends on the label of the interacting agents by differentiating the cases in the second sum.

Of course, the mass conservation property holds only for the whole system and not for each group. If we take $\varphi(v) = 1$ in (3.18) we can compute the evolution of the mass of agents with label i as

$$\frac{d}{dt} \rho_i(t) = \lambda \sum_{j=1}^n \int_V T(v; i|j) f_j(t, v) dv - \lambda \rho_i(t). \quad (3.19)$$

This approach connects the micro-level belief dynamics to the macro-level adoption curve: as more agents update their beliefs and find the product worthwhile, the adoption rate increases.

3.3.4 Time-scale separation

The description of the evolution of a system by means of the Boltzmann equation is tied to the microscopic quantities that form this equation. On the one hand, this means that it may not be best suited to understand the long-term behavior of the analyzed system. On the other hand, if two events occur on very different time scales, separated dynamics may better capture the phenomenon's essence. In this model, agents revise their beliefs about the product's value relatively frequently, whereas actual adoption decisions – once beliefs cross a threshold – occur less often. Capturing this time-scale separation is essential for modeling innovation diffusion in a realistic and tractable way.

We address this by distinguishing between two types of dynamics:

- Belief dynamics: how agents update their perceived value of the product through interaction and information flow.
- Adoption dynamics: how agents decide to switch from non-user to user status once their belief crosses a threshold.

Formally, let us consider a system where two types of interactions occur: frequent updates to valuation (e.g., social influence, advertising), and infrequent decisions to adopt. Each of these processes contributes differently to the overall evolution of the system. Let $f(t, v)$ be the distribution of agents over valuations at time t . The full model includes two terms: one governing fast (frequent) interactions and another governing slow (infrequent) transitions:

$$\frac{d}{dt}f(t, v) = \lambda Q(f, f)(t, v) + \mu P(f, f)(t, v), \quad (3.20)$$

where $Q(f, f)$ describes belief updates, $P(f, f)$ describes label transitions such as adoption and $\lambda \sim 1$ and $\mu \sim \delta^{-1}$, $\delta \rightarrow 0$ are the respective frequencies, representing time-scale separation. Furthermore, let us define a new time scale $\tau := \delta^{-1}t$ and refer to it as the fast time. The slow time will be t . We can now re-scale the distribution function as

$$\tilde{f}(\tau, v) = f(t, v). \quad (3.21)$$

Using this definition, we can split the equation (3.20) in the two scales, given that the contribution of $P(f, f)$ is of the same order as δ , as

$$\begin{aligned} D_\tau f(\tau, v) &= P(f, f)(\tau, v) \\ D_t f(t, v) &= Q(f, f)(t, v). \end{aligned} \quad (3.22)$$

The interpretation of this is the following. The system evolves by the action of the fast process $P(f, f)$ on the fast time τ toward its possible equilibrium configuration, undisturbed by the slow process $Q(f, f)$. On the slow time scale t this configuration evolves by action of the slow process.

3.3.5 Quasi-invariant limit

It is difficult to use the Boltzmann equation to analytically evaluate the equilibrium distribution of an observable. The quasi-invariant interactions limit allows us to do

so. If we assume that each binary interaction produces a small shift in the value of the perceived value v and we re-scale the time in such a way that one unit contains many of those interactions, we can study how the system behaves asymptotically. Mathematically, we re-write the binary interaction as

$$v' = v + \varepsilon I(v, v_*) + \eta, \quad (3.23)$$

with η random variable of null mean and variance σ^2 and we assume $\varepsilon, \sigma^2 \rightarrow 0^+$. We then define a new time scale $\tau := \varepsilon t$ and a re-scaled distribution function $g(\tau, v) := f(\frac{\tau}{\varepsilon}, v)$. Boltzmann equation now reads:

$$\frac{d}{d\tau} \int_V \varphi(v) g(\tau, v) dv = \frac{1}{\varepsilon} \langle \iint_{V \times V} (\varphi(v') - \varphi(v)) g(\tau, v) \cdot g(\tau, v_*) dv dv_* \rangle. \quad (3.24)$$

Quantity σ^2 is related to the contribution to the shift in the microscopic value of the stochastic fluctuations due to self-reasoning, personal preferences of the agents and other unaddressed influences. Quantity ε is related to the contribution to the shift in the microscopic value of the interactions with the outside world. Let us call the limit value $\frac{\sigma^2}{\varepsilon} \rightarrow \gamma > 0$ as $\varepsilon, \sigma^2 \rightarrow 0^+$. If we take the series expansion and we perform the limit for $\varepsilon, \sigma^2 \rightarrow 0^+$ it is possible to write the so-called Fokker-Planck equation (Risken, 1996):

$$\partial_\tau g(\tau, v) = -\partial_v \left(g(\tau, v) \int_V I(v, v_*) g(\tau, v_*) dv_* \right) + \frac{\gamma}{2} \partial_v^2 g(\tau, v). \quad (3.25)$$

We will use this equation to investigate the asymptotic behavior of the distribution of the perceived value of the new product in the population.

3.4 Multi-agent Bass model

Let us assume the market consists of N agents. Each agent is either a *non-user* or a *user*. We encode this status with a label $i \in \{1, 2\}$ where $i = 1$ corresponds to *non-user* and $i = 2$ to *user*. Every agent is also equipped with a microscopic state $v \in V \subseteq \mathbb{R}_+$ that represents its perceived value of the product. We employ the weak Boltzmann equation (3.18) for the label-restricted distribution function $f_i(t, v)$ to study the evolution of the mass and mean perceived value in the population. We

report the equation here for clarity:

$$\begin{aligned} \frac{d}{dt} \int_V \varphi(v) f_i(t, v) dv &= \lambda \int_V \varphi(v) \left(\sum_{j=1}^n T(t, v; i|j) f_j(t, v) - f_i(t, v) \right) dv \\ &+ \mu \sum_{j=1}^n \iiint_{V^3} \langle \varphi(v'_i) - \varphi(v) \rangle f_i(t, v) f_j(t, v_*) G_i(\hat{v}) dv dv_* d\hat{v}. \end{aligned} \quad (3.26)$$

The first term in the right-hand side of the equation accounts for the changes in evaluation in group i due to purchases; the second term accounts for changes in the evaluation due to interactions between agents, i.e. word-of-mouth, or with a background $G_i(\hat{v})$, i.e. advertisement for *non-users* and actual usage of the product for *users*. Let us first focus on the evolution of group sizes. From equation (3.19) we have:

$$\frac{d}{dt} \rho_i(t) = \lambda \sum_{j=1}^N \int_V T(t, v; i|j) f_j(t, v) dv - \lambda \rho_i(t) \quad i = 1, 2. \quad (3.27)$$

In this dissertation, we do not allow agents to abandon the innovation. The transition probability matrix thus becomes:

$$\begin{aligned} T(v; 1|1) &= 1 - H(v - P) & T(v; 2|1) &= H(v - P) \\ T(v; 1|2) &= 0 & T(v; 2|2) &= 1. \end{aligned} \quad (3.28)$$

Where P is the product price and $H(\cdot)$ is the Heaviside step function (Weisstein, 2002). The presence of the step function as a transition probability means that the agents behave as rational entities and as soon as their perceived value matches the real price of the product and a transition event occurs, they become *user*. The evolution of the mass in the two groups is therefore:

$$\frac{d}{dt} \rho_2(t) = \lambda \int_P^{+\infty} f_1(t, v) dv = -\frac{d}{dt} \rho_1(t). \quad (3.29)$$

This functional form of the differential equation governing the diffusion is different from those appearing in Young (2009), where contagion, social influence and social learning are considered. Given the presence of the integral over the perceived values, we can think of this model as an opinion-dynamics based model for innovation diffusion.

Next, we specify the interaction rules used in the second term of the right-hand side of equation (3.26). These rules differ for users and non-users:

- Non-users update their valuation based on word-of-mouth and external information (e.g., advertising):

$$v'_{1,i} = v - a_i(v - v_*) - b(v - \hat{v}) + v\eta. \quad (3.30)$$

a_i captures the strength of peer influence (word-of-mouth), and b captures the effect of external signals (advertising). Parameters a_i and b can be stochastic but the mean-field approach of equation (3.9) will average out any distributional effect. The structure of the interaction allows for different word-of-mouth parameters depending on the group of the interacting agent. This models, for instance, the situation where agents value more the opinion of a direct user rather than the opinion of a non-user. The stochastic term η is a random variable with $\langle \eta \rangle = 0$ and variance σ^2 . It accounts for unobserved mechanisms that influence the perceived value, such as personal preferences, internal reasoning or external events. Here, v_* represents the perceived value of the interacting agent and \hat{v} the advertisement and is sampled from $G_1(\hat{v})$.

- Users revise their valuation based solely on usage experience and product quality:

$$v'_2 = v - c(v - \hat{v}) + v\xi. \quad (3.31)$$

c is the experience-driven adjustment parameter, and ξ is the stochastic term (random variable with $\langle \xi \rangle = 0$ and variance σ^2). Here, \hat{v} represents the actual value of the product and is sampled from $G_2(\hat{v})$.

The background $G_1(\hat{v})$ in this model reflects both the content of the advertising campaign and individual skepticism about it. For simplicity, we model it as a distribution with two point masses:

$$G_1(\hat{v}) = (1 - w)\delta(v) + w\delta(v - P), \quad (3.32)$$

where $\delta(\cdot)$ is the Dirac delta function, P is the product price, and $w \in [0, 1]$ represents the effectiveness or credibility of the campaign. Conversely, background $G_2(\hat{v})$ represents the actual value of the technology. We assume:

$$G_2(\hat{v}) = \delta(v - rP), \quad (3.33)$$

where $r > 0$ is a coefficient capturing the perceived quality or real value of the product relative to price. We need to ensure that the interactions do not bring the perceived value outside its domain, i.e. $v'_i \geq 0, i = 1, 2$. This holds for

$$\eta \geq a_i + b - 1 \quad (3.34)$$

$$\xi \geq c - 1. \quad (3.35)$$

In this chapter, we focus on regimes where a_i, b, c and σ^2 are sufficiently small to guarantee non-negativity. While this interaction rule may appear mechanical, it captures a compromise-based updating process similar to bounded rationality models in behavioral economics. Each agent updates its valuation by averaging with others, reflecting limited cognitive processing and social influence.

Let us focus on the *non-users* and assume that decision-making occurs on a slower timescale than information acquisition. Hence, agents make up their minds before buying the product and impulse purchase is ruled out. This time lag is realistic and reflects the cognitive and logistical delays between forming an opinion and acting on it. Under this assumption, we can employ the hydrodynamic limit described in subsection 3.3.4 and divide into two temporal scales the label switch process and the interaction process. If we take $\lambda = 1$ and $\mu = \delta^{-1}, \delta \rightarrow 0$, the weak form of the split Boltzmann equation (3.22) for the first group reads:

$$\frac{d}{d\tau} \int_V \varphi(v) \tilde{f}_1(\tau, v) dv = \sum_{i=1,2} \iiint_{V^3} \langle \varphi(v'_{1,i}) - \varphi(v) \rangle \tilde{f}_1(\tau, v) \tilde{f}_i(\tau, v_*) G_1(\hat{v}) dv dv_* d\hat{v} \quad (3.36)$$

$$\frac{d}{dt} \int_V \varphi(v) f_1(t, v) dv = - \int_V \varphi(v) H(v - P) f_1(t, v) dv \quad (3.37)$$

In the same way, we can write them for group 2. The interpretation is as follows: the distribution $\tilde{f}_i(\tau, v)$ will quickly evolve and reach equilibrium following the interaction dynamics. Once this equilibrium is reached, the slow adoption process updates the group composition, prompting a new belief adjustment cycle, and so on.

The evolution of mass and mean perceived value on the fast scale will read:

$$\frac{d}{d\tau}\rho_1(\tau) = \frac{d}{d\tau}\rho_2(\tau) = 0 \quad (3.38)$$

$$\frac{d}{d\tau}m_1(\tau) = -a_2\rho_2(m_1 - m_2) - b(m_1 - wP) \quad (3.39)$$

$$\frac{d}{d\tau}m_2(\tau) = -c(m_2 - rP). \quad (3.40)$$

Let us note here that the dynamics of the mean perceived value for the non-users m_1 does not depend on the intra-group word-of-mouth parameter a_1 . On the slow scale t , the evolution of the mass for the two groups is the same as in (3.29):

$$\frac{d}{dt}\rho_2(t) = \int_P^{+\infty} f_1(t, v)dv = -\frac{d}{dt}\rho_1(t), \quad (3.41)$$

where the distribution $f_1(t, v)$ is the equilibrium distribution reached by the fast process.

In order to evaluate the equilibrium distribution $f_1(t, v)$ we employ the Fokker-Planck equation introduced in subsection 3.3.5. The quasi-invariant limit is satisfied for a_i, b and $\sigma^2 \rightarrow 0^+$. We set $\gamma = \frac{\sigma^2}{b}$ and $\alpha_i = \frac{a_i}{b}$. The interpretation of this is that γ represents the ratio between the importance of the self-reasoning process and the interactions and α the ratio between the word-of-mouth effect and the advertisement effect. In this case, Fokker-Planck equation (3.25) reads:

$$\partial_\tau f_1(\tau, v) = -\partial_v \left[f_1(\tau, v) \left(\sum_{i=1,2} \iint_{V^2} [\alpha_i(v_* - v) + (\hat{v} - v)] \cdot f_i(\tau, v_*) G(\hat{v}) dv_* d\hat{v} \right) \right] + \frac{\gamma}{2} \partial_v^2 (v^2 \cdot f). \quad (3.42)$$

We are looking for stationary solutions. Hence, we impose the time derivative to 0 and integrate:

$$0 = \partial_v \left[-f_1^\infty(v) \cdot (\alpha_1\rho_1 m_1 + \alpha_2\rho_2 m_2 + wP - (1 + \alpha_1\rho_1 + \alpha_2\rho_2)v) + \frac{\gamma}{2} \partial_v (v^2 \cdot f_1^\infty(v)) \right]. \quad (3.43)$$

This holds if the left-hand side of the equation under the derivative is constant. Let us assume it to be null to allow $f_1^\infty(v) = 0$ to be a solution. We obtain, as a stationary

solution on the fast time scale:

$$f_1^\infty(v) = C \cdot \frac{e^{-\frac{2}{\gamma} \frac{\alpha_1 \rho_1 m_1 + \alpha_2 \rho_2 m_2 + wP}{v}}}{v^{2 + \frac{2(1 + \alpha_1 \rho_1 + \alpha_2 \rho_2)}{\gamma}}}. \quad (3.44)$$

Where $C = K\rho_1$ is the normalizing constant over ρ_1 . If, on the fast time, we solve equation (3.39) we obtain that the mean perceived value for group 1 is:

$$m_1(\tau) = C_0 \cdot e^{-(a_2 \rho_2 + b)t} + \frac{a_2 \rho_2 r P + b w P}{a_2 \rho_2 + b}. \quad (3.45)$$

This means that the value of m_1 will go to $\frac{a_2 \rho_2 r + b w}{a_2 \rho_2 + b} \cdot P$ as τ goes to infinity. Similarly, m_2 will go to rP .

Equation (3.44) is now fully defined: the steady state is reached with respect to the fast time τ and thus the curve is parameterized by macroscopic quantities $\rho_1(t)$, $\rho_2(t)$ and $m_2(t)$, that evolve on the slow time t . Quantity m_1 is expressed as a function of ρ_2 and m_2 . Quantity m_2 goes to rP . In conclusion and by re-integrating the market size N , we can write the equation for the slow time as:

$$\frac{d}{dt} \rho_2(t) = N(1 - \rho_2) K \int_P^{+\infty} \frac{e^{-\frac{2}{\gamma} \frac{P}{v} \left[\alpha_1 \rho_1 \frac{a_2 \rho_2 r + b w}{a_2 \rho_2 + b} + \alpha_2 \rho_2 r + w \right]}}{v^{2 + \frac{2(1 + \alpha_1 \rho_1 + \alpha_2 \rho_2)}{\gamma}}} dv. \quad (3.46)$$




Fig. 3.1 Example of perceived value distribution in the population.

The expression under the integral sign represents the equilibrium profile of the statistical distribution of the perceived value in the population. Figure 3.1 shows an example of this distribution. This profile exhibits a fat tail for high values of v . For clarity, let us recap here the quantities forming this equation:

- N is the size of the market.
- γ is linked to the variance of the distribution of perceived value among the non-users. It is a measure of the divisiveness of the product.
- P is the price of the product. It can vary over time to model different pricing strategies of the product.
- a_2 is the word-of-mouth parameter between groups. It affects how much each interaction with a user agent has the power to change the perceived value of non-users.
- b is the advertisement parameter. It affects how much each interaction with the advertisement campaign has the power to change the perceived value of agents.
- r is the real value of the product, as the mean value perceived by the users.
- w is the strength of the advertisement campaign. By spending more money or better designing the campaign, managers can modify this number over time. However, it is outside the scope of this dissertation to design a ready-to-use mapping tool to convert actual advertising effort into parameter w .
- α_i is the ratio $\frac{a_i}{b}$, with $i = 1, 2$.

Equation (3.46) describes the evolution of the *non-users* population in time. It states that the number of agents that become *user* in a given instant is equal to the number of agents whose perceived value *in equilibrium* is higher than the sale price of the product in that instant.

The kinetic approach proposed in this chapter shares some structural similarities with the mean-field game model introduced by Chaab et al. (2022), but also presents important conceptual and methodological differences. Both frameworks operate at the mean-field level, where each agent is influenced by the aggregate state of the population. However, while Chaab et al. adopt a mean-field game approach - modeling a Stackelberg game between a forward-looking firm and strategic consumers - our model follows a kinetic perspective rooted in statistical mechanics. In their framework, the firm (leader) commits to a dynamic pricing and advertising policy, while consumers (followers) are segmented into individualists, influenced mainly by

advertising, and conformists, who respond to social forces. Each consumer solves an optimization problem to determine the optimal time of adoption within a finite horizon. In contrast, the agents in our kinetic model are behaviorally simpler: they adopt the product as soon as their perceived value exceeds the price. This assumption reflects bounded rationality and informational frictions more typical of early-stage market behavior, where optimization over a full time horizon may be unrealistic. Heterogeneity, such as the distinction between individualists and conformists, can be incorporated in the kinetic model by assigning different interaction parameters (e.g., different coefficients a_i and b) to different agent classes. However, doing so may complicate closed-form solutions. One of the main advantages of the kinetic approach lies in its ability to yield an explicit adoption curve derived from micro-level behavioral rules. Unlike MFG-based models, which require numerical solutions, the kinetic model provides a continuous-time diffusion law in closed form. This makes it particularly useful for forecasting and estimation in data-sparse environments.

This chapter develops a micro-founded diffusion model aimed at explaining aggregate adoption dynamics through a parsimonious behavioral mechanism. In doing so, it deliberately abstracts from the marketing literature that treats adoption as a function of managerial levers, controllable decisions typically organized as product, price, place (distribution), promotion and related launch and channel strategies (Kotler et al., 2014). As a consequence, the parameters estimated here should be interpreted primarily as reduced-form summaries of underlying behavioral heterogeneity and social interaction, conditional on an implicit and unmodeled marketing environment, rather than as direct prescriptions for marketing interventions. At the same time, a subset of parameters can be mapped to simplified policy proxies (e.g., the advertising term as a representation of promotional pressure), but only at a high level of abstraction and without the institutional detail emphasized in marketing practice.

If we solve the integral in equation (3.46), we obtain an operative equation for the evolution of the diffusion process:

$$\frac{d}{dt}\rho_2(t) = N(1 - \rho_2)\tilde{\gamma}\left(1 + \frac{2(1 + \alpha_1\rho_1 + \alpha_2\rho_2)}{\gamma}, \frac{2}{\gamma}\left[\alpha_1\rho_1\frac{a_2\rho_2r + bw}{a_2\rho_2 + b} + \alpha_2\rho_2r + w\right]\right). \quad (3.47)$$

Here, $\tilde{\gamma}(\cdot)$ is the regularized lower incomplete gamma function. Let us note here that the explicit price parameter P cancels out explicitly, but the adoption dynamics still

depend on marketing levers through the real value parameter r and the advertisement parameter w . Equation (3.47) is computationally tractable and produces an S-shaped adoption curve, a hallmark of innovation diffusion. While equation (3.47) is the most general form of the model, simpler variants can be used in practice. For example, in subsection 3.5.3, we consider the case where users and non-users exert equal influence on others ($a_1 = a_2$). Another tractable simplification assumes that only user interactions are effective ($a_1 = 0$). In the next section, we evaluate the empirical performance of the Kinetic Innovation Diffusion (KID) model using sales data from four real-world products.

3.5 Results

3.5.1 Nonlinear Least Squares

We now evaluate the performance of the Kinetic Innovation Diffusion (KID) model on real sales data for four benchmark consumer products: color televisions, air conditioners, clothes dryers, and freezers. These products are commonly used in the diffusion modeling literature and have appeared in several foundational studies (Bass et al., 1994; Cosguner and Seetharaman, 2022; Jiang et al., 2006). We use the data for these products contained in Cosguner and Seetharaman (2022), and we report for reference the fittings obtained in the same paper, the most recent work on this topic. Model fitting was conducted using a nonlinear least squares method applied to the general form of the KID model, implemented in Python. The emphasis of this section is on the out-of-sample forecasting performance of the model, rather than the quality of the in-sample fit. Given the KID model's larger parameter space, a lower sum of squared errors (SSE) is expected compared to the benchmark models. This expectation is confirmed in table 3.1, which reports the in-sample SSE from the best-fitting KID model alongside the SSE values reported by Cosguner and Seetharaman (2022), and highlights in bold the smallest errors. We report in table 3.2 the parameters that yield these best fits, for completeness.

To obtain optimal parameters, we used a Monte Carlo parameter space exploration, repeatedly fitting the model from randomly sampled initial conditions. This stochastic approach reduces the risk of converging to local minima and ensures a more robust fit.

	KID	BM	GBM	BGDM	BLDM
TV	130,740	996,421	402,514	457,051	457,387
AC	53,635	341,468	294,659	90,727	89,550
Dryer	76,886	212,777	210,770	129,485	129,227
Freezer	123,521	242,672	209,251	136,767	136,728

Table 3.1 The SSE for the four products of the compared models.

	a_1	a_2	b	γ	w	r	N
TV	3.49 1e-3	9.58 1e-2	7.49 1e-3	0.165	0.533	0.989	51775
AC	4.04 1e-2	1.46 1e-3	4.30 1e-5	1.296	0.929	0.966	21694
Dryer	7.38 1e-3	3.83 1e-2	4.11 1e-3	0.069	0.7636	0.931	28182
Freezer	4.32 1e-3	3.56 1e-3	1.60 1e-5	1.691	0.834	0.893	63746

Table 3.2 The parameters of the kinetic diffusion model for the four products.

It is important to stress that the parameter estimates reported here should not be interpreted as accurate representations of the underlying behavioral parameters. Instead, they serve to demonstrate the model’s fitting and forecasting capability. For a more accurate estimation of real-world mechanisms, ensemble methods should be preferred over point estimates.

3.5.2 Bayesian Inference

In order to obtain an ensemble parameter calibration for this model, we implemented a Metropolis-Hasting type Markov Chain Monte Carlo (MCMC) algorithm to perform Bayesian inference on the parameter set θ . At heart, Bayesian inference is a disciplined way to learn from data under uncertainty. It combines two ingredients: a likelihood, which encodes how plausible the observed data are under a given model and parameter vector, and a prior, which encodes what is known (or ruled out) before seeing the data. Bayes’ theorem turns these into a posterior – an updated distribution over parameters that quantifies what the data have taught us and how uncertain we remain. In complex, nonlinear models, this framework is especially useful: it lets us encode constraints and knowledge about the phenomenon, propagate them through the model, and obtain uncertainty-aware estimates and predictions rather than point fits that ignore identifiability and noise.

The theoretical foundation of this technique is the Bayes theorem (Bayes, 1958), that reads:

$$P(\theta|y) = \frac{P(y|\theta)}{P(y)}P(\theta). \quad (3.48)$$

Here, $P(\theta)$ is called the *prior distribution* over parameters $\theta \in \Theta$ and encodes the experimenter's knowledge and beliefs about the problem and appropriate values for θ , before the observation of the data. $P(y|\theta)$ is the *likelihood function* and represents the probability of obtaining the observed data y given parameters θ . $P(y)$ is the *marginal likelihood* that normalizes the posterior and becomes central when comparing models. Bayesian inference aims to estimate the entire posterior distribution $P(\theta|y)$, not just a single point from it. Doing so yields uncertainty quantification that is both accurate and meaningful, because it reflects the experimenter's prior beliefs and the information contained in the data y . Most commonly, we achieve that with Markov Chain Monte Carlo (MCMC), which constructs a Markov chain whose long-run distribution is the target posterior. When the likelihood function is not known, it can be approximated using a number of techniques (Turner and Van Zandt, 2012), in this case we talk about approximated Bayesian computation, or ABC. The Metropolis–Hastings (MH) algorithm (Hastings, 1970) used in this dissertation works following a simple random walk scheme with an acceptance-rejection rule at each step. The chain converges to the posterior and we can approximate expectations by ergodic averages of the draws.

Bayesian methods have recently become standard in complexity settings, and especially in agent-based modeling (Dyer et al., 2022; Grazzini et al., 2017), because they let us encode constraints on interpretable parameters and carry uncertainty through to predictions, while offering a principled basis for comparing specifications.

Operatively, we first used global variance-based sensitivity analysis (Sobol indices) to assess the influence of each structural KID parameter on the sum of squared errors between the model and the observed adoption trajectories. Within plausible parameter ranges, the interaction coefficients a_1 and a_2 displayed negligible first-order and total indices, indicating that the available sales data do not carry enough information to identify them. By contrast, the imitation parameter b , the divisiveness parameter γ , the advertising strength w , the real value r , and the market potential N all exhibited non-trivial total Sobol indices, with N emerging as the dominant driver of the fit. We therefore fixed (a_1, a_2) at the non-linear least squares estimates reported in subsection 3.5.1 and restricted Bayesian calibration to the five-dimensional

Algorithm 1: Approximated Metropolis-Hastings sampling scheme

Input: Prior distribution $P(\theta)$, proposal distribution $Q(\cdot|\theta)$, observed data y , initial set θ_0 , number of iterations n .

for $r=1, \dots, R$ **do**
 | Simulate $x^{(r)}$ given θ_0 ;
end

for $i=1, \dots, n$ **do**
 | Sample θ^* from $Q(\cdot|\theta_{i-1})$;
 for $r=1, \dots, R$ **do**
 | Simulate $\tilde{x}^{(r)}$ given θ^* ;
 end
 | Evaluate approximated likelihood $P_{\theta^*}(y)$ using $\tilde{x}^{(r)}|_{r=1}^R$;
 | Set $\theta_i \leftarrow \theta^*$, $P_{\theta_i}(y) \leftarrow P_{\theta^*}(y)$ and $x^{(r)} \leftarrow \tilde{x}^{(r)}$ (accept) with probability:
 | $\alpha = \min\left(1, \frac{P(\theta^*)Q(\theta_{i-1}|\theta^*)}{P(\theta_{i-1})Q(\theta^*|\theta_{i-1})}\right)$;
 | Otherwise, set $\theta_i \leftarrow \theta_{i-1}$ (reject);
end

parameter vector $\theta = (b, \gamma, w, r, N)$. Table 3.3 reports the Sobol indices for each product. The negative indices are to be intended as noise.

	TV		AC		Dryer		Freezer	
	S1	ST	S1	ST	S1	ST	S1	ST
a_1	0.005	0.022	0.001	0.020	0.001	0.025	-0.003	0.015
a_2	-0.008	0.029	-0.008	0.025	-0.007	0.025	-0.005	0.019
b	0.005	0.056	0.004	0.056	0.003	0.054	0.000	0.046
γ	0.052	0.154	0.048	0.156	0.047	0.154	0.040	0.135
w	0.048	0.149	0.039	0.134	0.037	0.132	0.024	0.100
r	0.036	0.121	0.025	0.109	0.024	0.110	0.019	0.098
N	0.662	0.779	0.694	0.800	0.698	0.805	0.728	0.830

Table 3.3 The Sobol analysis results for each product.

For each of the four benchmark products (TV, AC, dryer, freezer), we constructed weakly informative priors on θ centred at the corresponding NLS point estimates, using scaled Beta distributions for the bounded parameters and Gamma distributions for γ and N . This choice respects the structural bounds of the model, regularizes poorly identified directions, and encodes prior information only at the level of means and coarse relative dispersion. The resulting priors are proper and product-specific, but share a common structure across technologies.

Given these priors, we defined a Gaussian likelihood for cumulative adoption, with the observation variance estimated from the residuals at each parameter value. The posterior distribution of θ for each product was explored using a two-stage Metropolis–Hastings random-walk algorithm. A pilot chain with diagonal proposal covariance provided an empirical estimate of the posterior covariance, which was then used to define an efficient proposal for a long main run. The proposal scale was adaptively tuned during warmup to achieve an acceptance rate in the range recommended by MCMC theory for random-walk Metropolis, i.e. 0.234 (Von Toussaint, 2011). The resulting chains yield approximate samples from the full joint posterior of the KID parameters for each product, enabling both parameter uncertainty quantification and posterior predictive analysis.

The posterior summaries confirm that the data are informative about the market potential N and, to a lesser extent, about the imitation intensity b and the real value r , while the posteriors for γ and w retain significant spread for some products. Posterior predictive adoption curves closely track the observed diffusion trajectories and provide credible intervals for key managerial quantities such as time to peak sales, peak magnitude, and saturation levels. Compared to the analysis based on point estimates, this Bayesian treatment makes explicit the range of plausible parameter configurations and forecasts compatible with the data, and thereby moves from single best fits to ensemble-based inference. Figure 3.2 plots the diffusion curves for the four products obtained by this ensemble method. The posterior distributions for each product and each parameter are reported in appendix B.

In the next section, we shift focus from in-sample fitting to the managerial relevance of the KID model, particularly in terms of forecasting future adoption trends—arguably the most critical use case for practitioners.

3.5.3 Early Stage Forecasting Method

Forecasts on future sales, to have managerial value, need to be reliable as early as possible in the diffusion process. Even a very accurate prediction of the peak in the sales or the market potential is useless for managers and decision makers, if it becomes available when it is too late to properly program production, stocking and distribution of the product. Forecasting the final 20% of a diffusion curve, as done in some prior studies (Cosguner and Seetharaman, 2022, e.g.) provides little

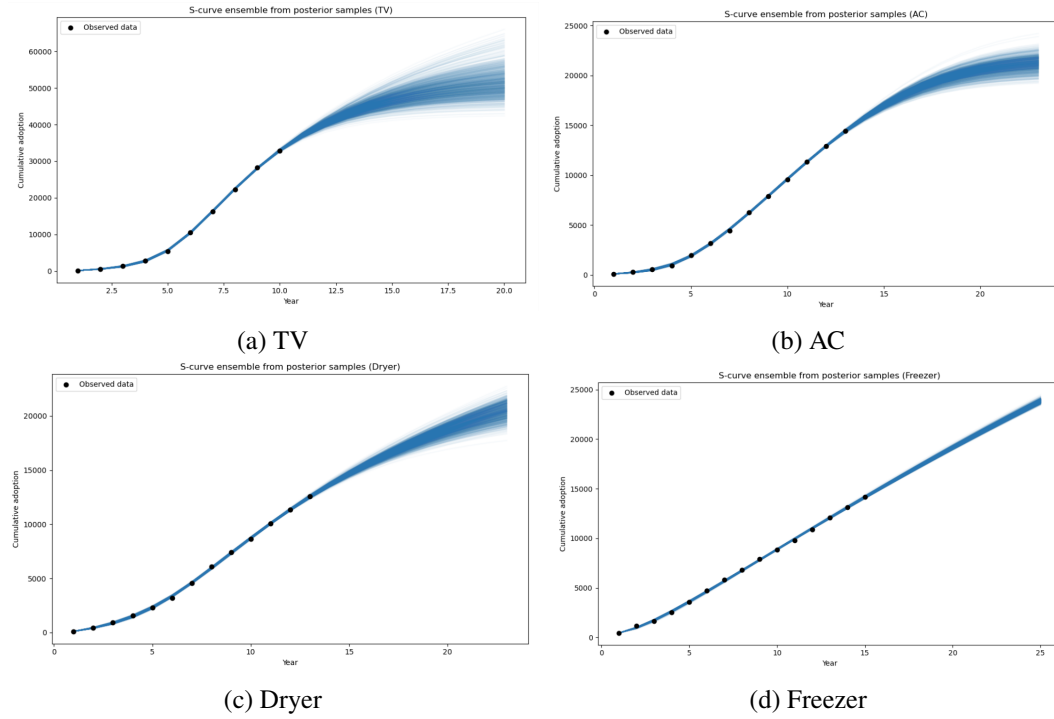


Fig. 3.2 Possible diffusion S-curves obtained by Bayesian inference of the parameters on the available data points (in black).

actionable insight for decision-makers. In practice, early-stage parameter estimation is often conducted through expert judgment or qualitative techniques (Lawrence and Lawton, 1981). The structure of the KID model supports this early estimation task. In particular, its theoretical formulation offers a direct expression, namely equation (3.46), for the distribution of perceived value in the population—even before the product has launched. At a pre-launch phase, where the number of users is zero and we have assumed $a_1 = a_2 = a$ and hence $\alpha_1 = \alpha_2 = \alpha$, that reads:

$$f_1^\infty(v) = C \cdot \frac{e^{-\frac{2}{\gamma} \frac{P \cdot w}{v} (1+\alpha)}}{v^{2 + \frac{2(1+\alpha)}{\gamma}}}. \quad (3.49)$$

This insight opens the door to parameter inference from pre-launch surveys: one can administer a survey to a target population exposed to the advertising campaign and fit equation (3.49) to the resulting distribution of perceived values. This procedure allows estimation of the advertising strength w and the ratio $\frac{1+\alpha}{\gamma}$. The market potential N is typically known or can be reasonably estimated by experienced marketing analysts. The real value r can be approximated using focus groups exposed to prod-

uct samples. Even if only one of the parameters a and b is known, the model remains usable by fixing one at a plausible level (e.g. $b = 0.001$), and computing the other via the relationship $a = b \cdot \alpha$. This framework allows practitioners to simulate alternative adoption trajectories under different strategic assumptions—making it a powerful tool for early-stage decision-making. We tested this hypothesis via a Monte Carlo implementation of this method. While actual belief data was unavailable for the benchmark products, we simulated the early-stage estimation process by perturbing the parameters obtained from the full-sample fit with random noise of $\pm 20\%$ of amplitude. These perturbed values were then used as initial estimates for fitting the model to only 3 or 5 early data points. We repeated this process 1000 times and computed mean and standard deviation of some significant quantities, namely market potential, time of the sales peak, sales peak and market saturation (calculated as the moment when 95% of the market potential is reached). These are key strategic quantities to plan production, stocking and distribution.

Pts	Prod.	Market Potential			Saturation		
		Mean	STD	Full S.	Mean	STD	Full S.
5	TV	63560	25946	53562	21	2	20
	AC	32820	11679	25319	25	4	30
	Dryer	16353	4945	15598	19	3	18
	Freezer	21240	7182	17059	22	3	20
3	TV	58610	13481	53562	22	4	20
	AC	38697	24940	25319	38	9	30
	Dryer	15205	2657	15598	16	4	18
	Freezer	32627	34652	17059	38	18	20

Table 3.4 Mean and standard deviation of forecast values of market potential and market saturation with 5 and 3 years of data compared with the values evaluated with the full sample, for the four products.

Tables 3.4 and 3.5 show the results for the considered products. Predictions made after the first 5 years are generally accurate within an acceptable statistical error (the only exceptions are freezer's peak sales and AC's saturation point). Predictions made after only 3 years are accurate for TV and dryer, while the standard deviation for AC and freezer is too large to derive accurate information. These findings demonstrate that the KID model is well suited for early-stage forecasting and can support managerial decisions at a time when such insights are most valuable. It offers quantitatively grounded predictions of the most critical planning parameters—without requiring the entire adoption trajectory to unfold.

Pts	Prod.	Peak			Flex		
		Mean	STD	Act.	Mean	STD	Act.
5	TV	7244	4316	5982	8	1	8
	AC	2925	1250	1828	9	1	8
	Dryer	1536	760	1523	10	2	8
	Freezer	1598	828	1205	12	3	13
3	TV	6493	2775	5982	8	1	8
	AC	2386	2446	1828	12	4	8
	Dryer	2308	1037	1523	8	2	8
	Freezer	680	225	1205	4	7	13

Table 3.5 Mean and standard deviation of forecast values of sales peak and flex point with 5 and 3 years of data compared with the actual values, for the four products.

3.6 Conclusions

This chapter has introduced a novel methodology for modeling new-product diffusion, grounded in the tools of statistical mechanics. The proposed Kinetic Innovation Diffusion (KID) model is a utility-driven, agent-based framework in which agents decide to adopt based on their individual perceptions of product value. These perceptions evolve through both peer interactions and exposure to advertising, with updates governed by a compromise-based rule. Unobserved individual variability is captured through stochastic noise. The model provides a closed-form analytical expression for the adoption curve, derived directly from micro-level behavioral assumptions. Importantly, it bridges the gap between aggregate diffusion models and simulation-based agent-based models. Unlike aggregate models, which rely on top-down assumptions about market dynamics, the KID model builds diffusion from the bottom up, treating it as an emergent phenomenon of individual belief dynamics. The results shown in section 3.5 confirm the validity of the model. The KID proved to explain better than the compared methods, namely BM, GBM, BLDM and BGDM, the sales data of four key products: color TV, air conditioning, clothes dryer and freezer. Additionally, an ensemble parameter calibration provided data-informed statistical distributions on the parameter set. Moreover, the structure of the KID model allows a new method for pre-launch or early-stage sales forecasts. This enables managers to use this model for strategic planning.

Conceptually, the KID model offers a synthesis between system dynamics and agent-based modeling. It retains the analytical tractability of aggregate models while

preserving the flexibility of micro-founded simulations. Features such as consumer segmentation, product abandonment, leader-follower dynamics, or nonlinear interaction rules can be naturally incorporated within this framework without requiring assumptions about the aggregate adoption curve.

Nonetheless, the model's analytical richness comes at the cost of mathematical complexity. As additional behavioral features are introduced, the model may become less tractable and the notation more cumbersome – an inherent trade-off relative to purely simulation-based ABMs.

In the next chapter, we further develop this model by adding an explicit perceived value dynamic and a non-stationary technology performance level. The resulting system is used to provide a rigorous mathematical formulation of a diffusion framework that is largely used by innovation professionals but has been somewhat overlooked by academic literature: the Hype Cycle.

Chapter 4

The Hype Cycle as an Emergent Behavior

4.1 General Context

In the mid-1980s, “expert systems” promised to bring machine intelligence into factories, finance, and medicine. Research groups scaled rapidly, venture funding flowed, and software and hardware companies were founded. The ever-increasing computational power of personal computers delivered promising results to the companies that invested in AI, in the form of cut-costs in testing and manufacturing. The “Live Experts on a Floppy Disk” (Kupfer, 1987) were bringing into the market the future that the space shuttle and *Star Wars* brought in the imagination of the public. But soon, the rule-based architecture of expert systems crumpled onto itself. Brittleness, opacity and increasing maintenance costs broke the dream of human-level intelligence in machines and by 1987 an entire industry had collapsed. Thousands of jobs and millions of dollars were lost. Budgets were cut, labs rebranded, and the term “AI” became a liability in grant applications. With time, expert systems were embedded into conventional models and made their way back into the market, with calmer and more realistic expectations upon them. This was not the first time artificial intelligence related enthusiasm faded after a period of excitement: in the early 1970s, the Lighthill report (Lighthill, 1973) in the UK and DARPA’s frustration with the Speech Understanding Research in the US put a stop to the enthusiasm that

earlier research on perceptrons and machine translation had sparked (see Crevier, 1993, for a complete and compelling history of research in artificial intelligence).

Artificial intelligence is not the only field where oscillations like this were observed. In the history of innovation, countless new technologies underwent fluctuating trajectories in the exposure to the public and in the enthusiasm of the people. New products and ideas from different domains followed a tumultuous path from the cradle to the general success and adoption, or to oblivion. From fuel cells (Konrad et al., 2012) to e-commerce (Konrad, 2006), from gene therapy (Van Lente et al., 2013) to hydrogen (Alkemade and Suurs, 2012), many innovations shared the same causal structure in this oscillating path: (i) an exogenous trigger (a launch, a program, a media wave) initiates an expectations surge; (ii) boundedly rational actors update on noisy signals and local feedbacks and build up expectations; (iii) real performance adjusts on slower, constrained timescales: inflated expectations are disappointed by reality; (iv) the correction produces negative effects that can impair capability accumulation; (v) if the innovation proves valid, a new and better-informed wave of attention rises driven by the actual development of the technology. This pattern is often looked at through the lenses of the Hype Cycle framework (Linden et al., 2003), introduced by the technology research and consulting firm Gartner Inc. in 1995.

Despite the many appearances of this pattern and the possibly material consequences it carries, we lack a formal, mathematically rigorous model that considers its complex nature. In this chapter, we show how the kinetic approach based on statistical mechanics presented in chapter 3 can fill this gap, providing a micro-founded mechanism – and a tractable dynamical system – for the emergence of the hype-disappointment-stabilization pattern, among other possible behaviors. This deeply non-linear path in the expectations of the public towards an innovation has its roots in the intertwined dynamics of different aspects of the introduction of a new product in a market: the adoption curve, the technology performance curve and the expectations curve itself. Each one of these phenomena is extensively studied in its own literature, but is the interplay of the three that produces the oscillating dynamic suffered by, for one, expert systems in the 1980s. The model allows for an exogenous initial media impulse but shows how endogenous amplification and correction emerge from the joint dynamics of perceived value (expectations), realized value (performance), and adoption. This results in a tractable dynamical system that is then studied with analytical and numerical methods to gain insights on the

determinants of peculiar behavior such as hype cycles, hype-disappointment curves, multiple peaks, or other observed patterns.

The chapter is organized as follows. The next section describes the Hype Cycle framework as introduced by Gartner Inc. and analyzes its criticalities and possibilities as expressed in the related academic literature. Section 4.3 builds on the model from the previous chapter by adding an explicit perceived value dynamic and a technology performance curve. Section 4.4 presents the analysis of the dynamical system that the model provides and studies its behavior, showing some numerical simulations. Section 4.5 concludes the chapter and discusses future directions.

4.2 The Hype Cycle Framework

Gartner's Hype Cycle (HC for short) is an influential framework that describes the early stages of technology deployment in terms of *expectations* or *visibility* to the public. By visualizing the phases through which the enthusiasm toward a new technology goes, it serves as a decision heuristic for companies to plan the engagement with the innovation (Fenn and Raskino, 2008). Gartner defines five sequential stages on the HC, from the initial event that sparks the enthusiasm to relative maturity of the technology, when adoption is around 20% to 30%:

1. *Technology trigger*: some public event generates public interest in some emergent technology. Typically, no commercial product exists at this stage, early investors might provide funding.
2. *Peak of inflated expectations*: the first generation of products come out and media interest increases, the still unclear potentialities of the new technology drive up the expectations of the public. As the number of users grows, the first problems with first-generation products become visible. Negative reviews start pulling enthusiasm down.
3. *Trough of disillusionment*: the technology does not live up to the inflated expectations and is discredited, media lose interest. However, the product improves thanks to early feedback and some users find useful applications.




Fig. 4.1 The Hype Cycle formed by the hype curve and the technological S-curve and the five phases. Adapted from Fenn and Time (2007).

4. *Slope of enlightenment*: experimentation and real-world experience lead to a better understanding of the technology, better products are developed and a more realistic wave of interest rises.
5. *Plateau of productivity*: the benefits of the technology are accepted and mainstream adoption begins.

The HC is obtained as a superposition of two effects: a human-centric bell-shaped hype curve, product of attraction to novelty, social contagion and heuristic decision making; and a technological S-curve that describes the performance evolution of the product (Fenn and Raskino, 2008). Figure 4.1 shows this superposition of curves and the five phases of the HC as explained by Gartner. The speed at which technologies go through the HC varies from one case to the other, and Gartner publishes every year a report where it positions developing technologies on the curve, assessing the estimated time for each technology to reach the plateau of productivity.

Despite practitioners have been using the Hype Cycle model for almost thirty years, academic literature has highlighted some criticalities that limit its applicability and repeatability. At a closer look at figure 4.1, two of such criticalities clearly appear. First, while the x-axis represents time, it is not well-defined what is represented on the y-axis; and second, it is not clear how can one perform a superposition, or a sum, of two quantities such as *hype* and *technological performance*, that differ in nature and, even if well-defined, surely would differ in units of measure (Steinert and Leifer, 2010). Fenn and Raskino (2008) ambiguously refer to what lays on the y-axis both as *expectations* and *visibility*, without giving an operational definition. This has

led different academic studies to use different concepts and proxies to rigorously assess technological hype cycles. The number of articles appeared on mainstream or technical newspaper, both with (Konrad et al., 2012) and without (Alkemade and Suurs, 2012; Jarvenpaa and Makinen, 2008a,b; Konrad, 2006; Van Lente et al., 2013) sentiment assessment, the number of academic papers published (Campani and Vaglio, 2015; Rachul and Zarzeczny, 2012), the number of new patents granted (Dehghanimadvar et al., 2022; Khodayari and Aslani, 2018), the amount of search traffic registered (Jun, 2012a,b; Jun et al., 2014), and the count of comments posted on social media (de Cuveland et al., 2025) all appear in the literature as proxy for *expectations* or *visibility*. Additionally, Dedehayir and Steinert (2016) observe several empirical incongruences in the positioning of the innovations on the curve in Gartner reports: excessive duration, multi-year peaks, entries at plateau and removals before completion. These criticalities result in a limited repeatability of the results displayed by Gartner, and in an overall non-scientific method that yields little forecasting reliability.

In the model described in chapter 3, the perceived value v serves us as well-defined *expectations* and we show how, under certain conditions, its trajectory in time draws the shape typical of the Hype Cycle, without the need of factitious sums of different quantities. This novel approach differs from the other – rare – works that attempt to formalize the microscopic dynamic beneath hype cycle formation, but also shares some similarities. Silvestrini et al. (2017) use a method from quantum physics, namely rate equation, to describe the movements of agents in a population on levels of interest towards a technology. They define *visibility* as the the number of new users per time unit and show parameter combinations that yield hype cycle curves. Hashemi et al. (2018) focus on the microscopic mechanism of hype creation. They propose a multi-agent model that shows how rational interacting *technology experts* can produce *social bubbles* that reads like hype cycle dynamics. The interplay between hype, adoption and performance is investigated by Shi et al. (2023). Here, authors build an agent-based model for green energy technology diffusion that take into account these three aspects, and show simulations where hyped dynamics appear.

4.3 Model

We have argued that Hype Cycles emerge when expectations adjust faster than realizations, and that their foundations lie in the interplay between the social sphere – opinion dynamics and adoption – and the technological sphere – performance trajectories. Hence, people and technology are the main components of this model. Let us recap here the kinetic model for innovation diffusion described in the previous chapter and add an explicit perceived value dynamics to it. Each agent in the model is equipped with a positive perceived value of an innovation $v(t) \in \mathbb{R}_+$ and with a binary adoption status $i(t) \in \{1, 2\}$, where $i = 1$ denotes agents who have not bought the technology, or *non-users*, and $i = 2$ denotes agents who have, or *users*. State variable $v(t)$ will change in time through interactions between agents and with the media and possibly self-reasoning; while state variable $i(t)$ will change according to a purchase mechanism. On the other hand, the technology is characterized by a positive price $P(t) \in \mathbb{R}_+$ that we now take non-stationary, and by a real value $P(t) \cdot r(t)$, being $r(t)$ the real value coefficient, also non-stationary. The media campaign that can influence agents' perception is defined by two parameters: the frequency with which agents come into contact with it, $\mu_{media}(t)$; and the value conveyed by it, $P(t) \cdot w(t)$ – here expressed for convenience in relation to the price of the technology $P(t)$. We define $f_i(t, v)$ as the distribution function of agents with adoption status i over perceived values v at time t . This allows us to calculate the number $\rho_i(t)$ and the mean perceived value $m_i(t)$ of agents with adoption status i by integration. Being $\rho(t) = 1 \forall t$ and $m(t)$ the respective quantities in the whole population:

$$\rho(t) = \int_{\mathbb{R}_+} f(t, v) dv = \int_{\mathbb{R}_+} [f_1(t, v) + f_2(t, v)] dv = \rho_1(t) + \rho_2(t), \quad (4.1)$$

$$m(t) = \int_{\mathbb{R}_+} v \cdot [f_1(t, v) + f_2(t, v)] dv = \rho_1(t)m_1(t) + \rho_2(t)m_2(t). \quad (4.2)$$

Quantity $m_1(t)$ denotes the mean perceived value of the innovation among the *non-users* at time t and represents in this model the *expectations* that lie on the y-axis in the HC. Quantity $\rho_2(t)$ denotes the number of agents who have adopted the innovation at time t and constitutes the adoption process. Although both *users* and *non-users* update their perceived value over time, we define the hype cycle using the *non-users* mean perceived value because it is the variable that most closely corresponds to the Gartner narrative of market-wide pre-adoption expectations.

Agents interact with their environment and update their perceived value according to the set of rules described in chapter 3. This is a compromise-based rule that updates the perceived value of the agents with specific probabilities (that become frequencies when considering a large number of agents) and following different mechanisms according to the adoption status of each agent. Table 4.1 summarizes the update rules for clarity.

Table 4.1 Perceived value updates ($v \mapsto v'$) by group and channel.

Adoption Status	Type	Update rule	Frequency
Non-users	Agent-agent interaction	$v' = v - a_i (v - v_*)$	1
	Agent-media interaction	$v' = v - b (v - P \cdot w(t))$	$\mu_{media}(t)$
	Self-reasoning process	$v' = v + v \cdot \eta$	1
Users	Usage	$v' = v - c (v - P \cdot r(t))$	1
	Self-reasoning	$v' = v + v \cdot \xi$	1

v_* is the interacting agent's perceived value. Coefficients $a_i, b, c \geq 0$.

Agents that think the technology is valuable can make a purchase and change their adoption status from *non-user* to *user*. Purchase events occur with frequency λ and finalize if the perceived value of the interested agent $v(t)$ at time t is higher than the price $P(t)$ of the technology.

Together, the update rules and the purchase mechanism define a twofold dynamic where perceived values in a population of *non-users* initially distribute around the value conveyed by the media (**agent-media interaction**), with a variance that depends on the variance of η (**self-reasoning**). As adoption ρ_2 grows (**purchases**), the *users* start to communicate (**agent-agent interaction**) the real value they experience (**usage**) and, given their higher credibility as first-hand testers of the technology ($a_2 > a_1, b$), move the mean of the distribution of the perceived value among the *non-users* m_1 towards the real value $r(t) \cdot P(t)$. If this is high enough, the adoption process is successful and the *users* population continues to grow following an S-shaped curve until it reaches saturation, as it is described in chapter 3. Let us re-write for clarity the results of the previous chapter: the time-scale separation described in 3.3.4 yields, for the updating mechanism happening on the fast time, equations 3.38,

3.39 and 3.40. The solution of those yields:

$$\frac{d}{d\tau}\rho_2(\tau) = \frac{d}{d\tau}\rho_1(\tau) = 0, \quad (4.3)$$

$$m_1(\tau) = \frac{a_2\rho_2r + \mu_{media}bw}{a_2\rho_2 + \mu_{media}b} \cdot P + Ke^{-(a_2\rho_2 + \mu_{media}b)\tau}, \quad (4.4)$$

$$m_2(\tau) = rP + K'e^{-c\tau}. \quad (4.5)$$

On the slow time scale, the solution of equation 3.36 reads:

$$\frac{d}{dt}\rho_2(t) = -\frac{d}{dt}\rho_1(t) = \lambda \int_P^{+\infty} \bar{f}_1(t, v) dv, \quad (4.6)$$

$$m_1(t, \rho_2) = \frac{a_2\rho_2r(t) + \mu_{media}(t)bw(t)}{a_2\rho_2 + \mu_{media}(t)b} \cdot P(t), \quad (4.7)$$

$$m_2(t) = r(t) \cdot P(t). \quad (4.8)$$

Where $f_1(t, v)$ is the equilibrium distribution that results from the fast process. We have defined at this point the intertwined dynamics of adoption and expectations in the population, let us now discuss how to plug the technology performance curve into the system.

A large empirical literature characterizes technological improvement either as a time trend or as a experience effect. The *Moore's law* treats performance or cost as improving exponentially in time. *Wright's law* (Wright, 1936) states that unit cost falls as a power law of cumulative production. In a seminal reconciliation, Sahal (1979) showed that when cumulative output itself grows approximately exponentially in time, the two views become observationally hard to distinguish: *Wright's law* in cumulative production space induces a Moore-like exponential decline in time, and vice-versa. Modern tests on multi-technology panels (Nagy et al., 2013) confirm the practical equivalence under rapid output growth and quantify forecast errors. This equivalence does not imply unbounded improvement. Both relationships are local regularities useful to make predictions (Lafond et al., 2018): they say nothing about ultimate limits or design ceilings; S-curve models (Foster, 1986) take these limits into consideration, introducing a saturation effect in the growth as they are approached. They depict a slow takeoff, rapid gains as a design is exploited, deceleration as constraints appear, and possible succession to a new S-curve when an architectural change occurs (Dosi, 1982). The exponential trends effectively pool together successive S-curves across paradigms, even as each individual trajectory

saturates. According to Fenn and Raskino (2008) and to figure 4.1, we model realized performance r with a technology S-curve. This representation captures diminishing returns and a practical upper bound in performance. Technology S-curves are often plotted with performance against time, but this is accurate only if the amount of effort is relatively constant over time. "If the amount of effort invested in a technology decreases or increases over time, the resulting curve could appear to flatten out much more quickly, or not to flatten out at all" (Schilling and Esmundo, 2009). It is sensible then to use R&D cumulative investment as independent variable in the definition of r . In this paper, the focus is on how the interplay of different social and technological dynamics can produce the observed pattern of HC, and a precise empirical evaluation of the investment strategies of firms and of the causal relation between investment and growth is out of scope. Consequently, we model the cumulative investment as a function of adoption (cumulative sales) and cumulative expectations, with a exogenous positive investment term that account for unmodeled adjustments. Realized performance follows a sigmoid function of this investment. Mathematically:

$$r(I) = \frac{r_{\infty}}{1 + e^{-v[I-I_0]}} \quad (4.9)$$

$$I = k_1 \int_0^t m_1(s) ds + k_2 \rho_2 + \bar{I}(t). \quad (4.10)$$

Here, r_{∞} is the maximum performance, v the growth rate and I_0 the mid-point of the sigmoid; k_1 and k_2 are weights and $\bar{I}(t)$ is the exogenous investment.

We have now all the ingredients that we need to write the complete dynamical system for the three curves:

$$\begin{cases} m_1(\rho_2, r) = \frac{a_2 \rho_2 \cdot \mathbf{r} + b \mu_{media}(t) \cdot w(t)}{a_2 \rho_2 + b \mu_{media}(t)} \cdot P(t) \\ \frac{d}{dt} \rho_2(t) = \lambda \int_{P(t)}^{+\infty} f_1(t, v) dv \\ r(\rho_2, m_1) = \frac{r_{\infty}}{1 + e^{-v[k_1 \int_0^t \mathbf{m}_1(s) ds + k_2 \rho_2 + \bar{I}(t) - I_0]}} \end{cases} \quad (4.11)$$

With $\mu_{media}(t)$, $w(t)$ and $P(t)$ exogenous functions. System 4.11 lives on the slow time scale and its meaning is the following. Mean perceived value m_1 instantaneously collapses on the τ -asymptotic value parametrized by the slow moving variables ρ_2 and r , and by the exogenous parameters. The evolution of the mass of the adopters

ρ_2 is regulated by the purchase mechanism and depends on the fraction of *non-users* whose perceived value is higher than the price at time t . Distribution $f_1(t, v)$ depends on the microscopic updating rules and can be explicitly computed as an equilibrium distribution on the fast time scale. We will distinguish two cases: a deterministic case where stochastic self-reasoning is tuned down and the distribution $f_1(t, v)$ collapses into a Dirac delta on m_1 ; a noisy case with self-reasoning where distribution $f_1(t, v)$ is calculated via a Fokker-Planck equation (see subsection 3.3.5). Realized performance r grows with the investment that is proportional to cumulative mean perceived value and adoption. Price P and media/advertising parameters μ_{media} and w are taken as exogenous signals. In the next section, we will show what conditions are to be satisfied for the Hype Cycle or other patterns to emerge.

4.4 Analytical Results and Simulations

We analyze the system in a deterministic case, where the stochastic self-reasoning update rule is turned off. This generates a simplified dynamics where the conditions for the emergence of interesting patterns can be easily derived analytically. The cost for this is a piece-wise exponential adoption, different from the canonical S-curve. A smoothed dynamics with S-shaped adoption is retrieved when self-reasoning is turned back on, and agents behave with some individuality.

4.4.1 Deterministic case

Turning self-reasoning off results in agents all behaving deterministically in the same way, and hence the distribution function of the perceived value in the population collapses to the Dirac delta $f_1(v) = \delta(v - m_1)$. The second term in 4.11 reduces to:

$$\dot{\rho}_2(t) = \lambda(1 - \rho_2)\mathbb{1}[m_1(t) \geq P(t)]. \quad (4.12)$$

This yields an adoption process that is piecewise exponential under a hard threshold. Adoption abruptly stalls when $m_1 \leq P$ and is purely exponential when $m_1 \geq P$. This is consistent with the inertial case discussed by Young (2009). The deterministic assumption does not directly affect the evolution of m_1 or r . We have defined a HC as a peak in expectations followed by a slow recovery later on, so it comes

natural to study the first derivative of m_1 to find conditions on the formation of hyped-dynamics. We can calculate the total derivative with respect to time as sum of partial contributions:

$$\dot{m}_1 = \frac{\partial m_1}{\partial P} \dot{P} + \frac{\partial m_1}{\partial \mu_{media}} \dot{\mu}_{media} + \frac{\partial m_1}{\partial w} \dot{w} + \frac{\partial m_1}{\partial r} \dot{r} + \frac{\partial m_1}{\partial \rho_2} \dot{\rho}_2 \quad (4.13)$$

Our interest lies in the relationship between realized performance, conveyed value, and price in relative terms. Accordingly, we normalize the price to $P \equiv 1$ and analyze the dynamics in terms of μ_{media} , w , r and ρ_2 . To compute each partial derivative, it is useful to define a positive quantity $D = a_2 \rho_2 + b \mu_{media}$, that represents the denominator in the definition of m_1 . It is straightforward to obtain:

$$\frac{\partial m_1}{\partial \mu_{media}} = \frac{a_2 b \rho_2 (w - r)}{D^2}, \text{ higher media exposition raises } m_1 \text{ if } w > r. \quad (4.14)$$

$$\frac{\partial m_1}{\partial w} = \frac{b \mu_{media}}{D}, \text{ stronger media content raises } m_1. \quad (4.15)$$

$$\frac{\partial m_1}{\partial r} = \frac{a_2 \rho_2}{D}, \text{ higher realized performance raises } m_1. \quad (4.16)$$

$$\frac{\partial m_1}{\partial \rho_2} = \frac{a_2 b \mu_{media} (r - w)}{D^2}, \text{ higher adoption lowers } m_1 \text{ if } w > r. \quad (4.17)$$

We now isolate the mechanism producing a peak–trough–recovery in non-user expectations m_1 . Under monotone adoption and monotone performance, the cycle arises from a competition between media growth, adoption growth, and performance catch-up. At the beginning of the cycle $\rho_2 \approx 0$ and so $m_1 \approx P \cdot w$. A sufficient trigger is simply $w(0^+) > r(0^+)$, that reflects an initial overestimation of the capabilities of the technology by the media. The first part of the dynamic is dominated by the update of the perceived value on the fast time scale that exponentially collapses to $P \cdot w(0^+)$. If $w(0^+) > 1$, this sparks adoption and the evolution of the system on the slow time scale. The majority of the studies that aim to empirically assess hype cycles through media exposure (see section 4.2 for references) focus on the evolution of number of appearances of the technology in articles on mainstream or technical newspapers. They agree that a hyped dynamic takes place when this count, that in our model can be reasonably linked to μ_{media} , first rises to the peak of inflated expectations and later falls down in the trough of disillusionment. We analyze the system assuming this type of unimodal behavior for $\mu_{media}(t)$ and a fixed value $w(t) = w(0^+) > 1, \forall t$. Let us note here that both adoption ρ_2 (exponential) and

realized performance r (sigmoid) are non-decreasing $\forall t$. With these assumptions, the evolution of mean perceived value reduces to

$$\dot{m}_1 \propto \frac{A}{D^2}(\rho_2 \dot{\mu}_{media} - \mu_{media} \dot{\rho}_2) + \frac{\rho_2}{D} \dot{r}, \quad A := b(w - r) > 0. \quad (4.18)$$

Re-arranging the equation in growth-rate form

$$\dot{m}_1 \propto \frac{A}{D} \mu_{media} \rho_2 (g_\mu - g_\rho) + \rho_2 \dot{r}, \quad (4.19)$$

with $g_\mu := \dot{\mu}_{media}/\mu_{media}$ and $g_\rho := \dot{\rho}_2/\rho_2$, makes the hype cycle mechanics transparent. Early hype rise occurs when the media-adoption race favors media growth: $g_\mu > g_\rho$, with \dot{r} initially small. The peak ($\dot{m}_1 = 0$) arrives when

$$g_\mu - g_\rho = -\frac{\dot{r} \cdot (a_2 \rho_2 + b \mu_{media})}{A \mu_{media}}, \quad (4.20)$$

so a larger gap $w - r$ postpones the peak, while faster performance catch-up \dot{r} brings it forward. Immediately after the peak, expectations enter the disillusionment phase whenever $\dot{m}_1 < 0$. Using the growth-rate form, disillusionment occurs when the media-adoption race turns sufficiently negative to overwhelm performance catch-up, i.e.

$$g_\mu - g_\rho < -\frac{\dot{r} \cdot (a_2 \rho_2 + b \mu_{media})}{A \mu_{media}}. \quad (4.21)$$

This refines the narrative condition $g_\mu < 0$ and $g_\rho > 0$: media decay and rapid adoption growth pull expectations down, but a trough is reached only if the negative drift generated by $(g_\mu - g_\rho)$ dominates the positive term \dot{r} .

The trough of disillusionment corresponds to the next zero-crossing of \dot{m}_1 after the peak (with $\ddot{m}_1 > 0$), which typically occurs once media influence has weakened and/or the gap $w - r$ has narrowed enough that the performance term can compensate the negative media-adoption drift. Finally, recovery follows as user-informed word-of-mouth overcomes media influence due to increased adoption and performance approaches claims. In particular, as $w - r \rightarrow 0$ (so that $A = b(w - r) \rightarrow 0$), the first term $\frac{A}{D} \mu_{media} (g_\mu - g_\rho)$ vanishes and the dynamics of m_1 becomes dominated by $\dot{r} \geq 0$, implying a return to growth and, in the long run, $m_1 \rightarrow r$, consistently with Gartner's description of the HC.

These conditions are general within the present setting, in the sense that they do not require explicit functional forms for $\rho_2(t)$ and $r(t)$, and they hold as long as adoption and realized performance are non-decreasing and $\mu_{\text{media}}(t)$ exhibits a unimodal profile.

4.4.2 Noisy case

When we re-introduce self-reasoning, agents no longer share the same perceived value and we find that the distribution of perceived value in the population is again equation 3.44:

$$\tilde{f}_1(v) = C \frac{e^{-\frac{2}{\gamma} \frac{P}{v} (\alpha_1 \rho_1 \frac{m_1}{P} + \alpha_2 \rho_2 r + w)}}{v^{2 + \frac{2(1 + \alpha_1 \rho_1 + \alpha_2 \rho_2)}{\gamma}}}, \quad (4.22)$$

where $\gamma > 0$ parametrizes self-reasoning variance, $\alpha_1 = a_1/b$, $\alpha_2 = a_2/b$ and $C > 0$ is a normalizing constant. The intuition is that idiosyncratic reflection and heterogeneous information make some people more optimistic and some more skeptical at any given time.

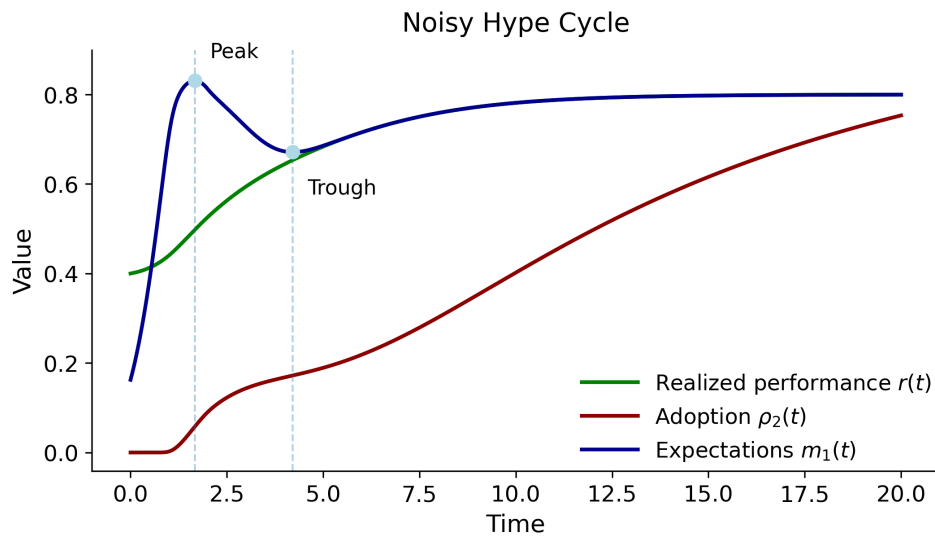


Fig. 4.2 The evolution of expectations, adoption and realized performance in the dynamical system.

This heterogeneity turns the hard purchase threshold into a smooth adoption hazard. Instead of adopting if and only if $m_1 \geq P$, the adoption rate is proportional

to the tail probability that an agent's perceived value exceeds the price. Formally,

$$\dot{\rho}_2 = (1 - \rho_2)\lambda \int_P^{+\infty} \tilde{f}_1(v) dv. \quad (4.23)$$

This tail behaves like a regularized incomplete gamma function: it rises smoothly with m_1 and falls with P , and depends continuously on adoption ρ_2 and realized performance r through the parameters of the distribution. Two immediate consequences follow: the adoption becomes broadly sigmoidal and the activation threshold relaxes, smoothening the dynamic.

Crucially, the hype cycle conditions derived in the previous subsection carry over unchanged in structure. The only change is that $\dot{\rho}_2$ now comes from the smooth hazard rather than an on/off switch, but retains its non-decreasing property. This is more realistic: even when m_1 is slightly below price P adoption still progresses rather than stalling abruptly. The deterministic case can be seen as a limit case of this broader form: as the self-reasoning variance parameter γ tends to zero, the perceived value distribution collapses to a point mass, the tail probability converges to an indicator function and the deterministic adoption law is recovered. Figure 4.2 shows a realization of the Hype Cycle with the complete (noisy) dynamic and time-varying μ_{media} . We omit a picture for the deterministic case as the hard threshold results in poorly informative and frankly quite ugly plots.

Heterogeneity at the fast scale smooths the adoption decision and delivers the empirically observed S-curve without changing the core mechanism behind hype–disappointment that emerged in the deterministic case: the initial overshoot is media-driven, with the increasing attention toward the technology pumping the hype, the correction arrives when media recedes faster than performance catches up, and the recovery as promised performance realizes. This is consistent with the findings of the majority of works that empirically assess technological hype cycles. In the next subsection, we will challenge this thesis by arguing that hyped dynamics can, in fact, emerge even if the exogenous media signal is fixed. The need for increasing media interest driving up the expectations derives from the fact that, in the time-scale separation approximation, perceived value adjusts instantly on the slow time. Using an agent-based simulation, we are able to relax the time-scale separation and retrieve a hype cycle dynamic with fixed w and μ_{media} .

4.4.3 Agent-based simulation

The analytical system is the right tool to extract clean conditions for the formation of HCs, but it rests on simplifying assumptions such as time-scale separation, mean-field closure for the perceived-value distribution, infinite N . A Direct Simulation Monte Carlo (DSMC) implementation lets us relax these: we can let interaction, self-reasoning, and adoption operate on comparable time-scales and we can expose finite- N stochasticity. In short, the simulation is a robustness check that the HC mechanism is not an artifact of the particular assumptions necessary to formulate a closed-form dynamical system.

We adapt the kinetic DSMC used for Boltzmann-type models to our two-population setting (non-adopters/adopters with perceived value v). The scheme follows the uncoupling principle: in a small step Δt , we apply interaction (“collisions”) and self-reasoning and label switching as two sub-moves. The idea of the simulation scheme (from Nanbu, 1980) is to update one particle at a time, sampling whether it collides and with whom; this preserves molecular chaos and is provably convergent in Babovsky’s variant (Babovsky and Illner, 1989).

The simulation and the dynamical system are equivalent: with the same parameters and with small Δt and $\mu \gg \lambda$, they produce the same curves for $m_1(t)$, $\rho_2(t)$ and $r(t)$ (see figure 4.3a). This validates that the aggregate model is the mean-field limit of the micro rules (Krook and Wu, 1977). The simulation allows us to relax the time-scale separation hypothesis by setting μ and λ on the same order of magnitude. Figure 4.3b shows the outcome of such run, where we also kept the media signal constant in time. A clear HC pattern is still visible with this assumptions: even with all the exogenous signals fixed, the Hype Cycle emerges as a product of the intertwined phenomena of early optimistic promises, adoption and delayed realized performance.

4.4.4 Figures Generation

The figures in this section are illustrative, not calibrated. Their purpose is to demonstrate that the proposed model can generate the canonical hype-disappointment-recovery pattern under economically interpretable conditions; they are not fits to any specific historical episode. We selected parameter values by guided tuning within

Algorithm 2: Nanbu-Babovsky scheme for Direct Simulation Monte Carlo

Input: Number of agents N , initial perceived values $V_0 = v_1^0, \dots, v_N^0$ and adoption states $X_0 = x_1^0, \dots, x_N^0$, time step Δt , number of steps n_{TOT} ;

for $t = 1, \dots, n_{TOT}$ **do**

Compute $\rho_2(t)$, $m_1(t)$ and $r(t)$;

repeat

Randomly select a pair of agents i and j ;

for $h = i, j$ **do**

Sample $\Theta \sim \text{Bernoulli}(\mu\Delta t)$;

if $\Theta = 1$ **then**

Update $v_h^t \leftarrow v_h^j$ according to adoption states x_i^{t-1} and x_j^{t-1} and update rules 3.30 and 3.31;

else

$v_h^t \leftarrow v_h^{t-1}$;

end

Sample $\Xi \sim \text{Bernoulli}(\lambda\Delta t)$;

if $\Xi = 1 \wedge x_h^{t-1} = 0 \wedge v_h^{t-1} \geq P$ **then**

$x_h^t \leftarrow 1$;

else

$x_h^t \leftarrow x_h^{t-1}$;

end

end

until no unused pairs are left;

end

plausible ranges reported in subsection 3.5.1. In particular, coefficients governing (i) media influence and content, (ii) interaction/usage coefficients, (iii) the purchase rate λ , and (iv) the performance map $r(\cdot)$ were adjusted to produce a single, clearly interpretable hype cycle with recovery. No attempt was made to match magnitudes, dates, or peak widths to data; all axes are normalized and shown in model time.

The media strength μ_{media} is modeled as a unimodal bump in the noisy case plot and held constant in the simulation, the media content $w(t)$ is kept above price at launch to trigger the hype, and $P(t)$ is normalized to 1 for readability. Initial conditions set low adoption and performance below claims; the realized-performance ceiling r_∞ is set to 0.8 in the noisy case plot and to 0.9 in the simulation.

Because these curves are not empirical fits, they should be read as existence proofs and mechanism demonstrations. A natural continuation is to assemble data on people perceived value via survey, media volume/content, adoption proxies, and

technology performance and then estimate the model (e.g., via MLE or Bayesian methods) and validate out of sample. That calibrated exercise falls outside the scope of this dissertation and is left for future work.

4.5 Conclusions

In this chapter, our goal was to deliver a rigorous, micro-founded model that explains hype-disappointment-recovery dynamics as an emerging property of intertwined social and technological phenomena. Existing Hype Cycle narratives are descriptive, conflate distinct processes (media attention vs. user experience vs. adoption), and offer no testable conditions for when a hype–disappointment–recovery sequence should emerge. We addressed these shortcomings by separating—and then coupling—three dynamics: (i) perceived value (expectations) generated through social interaction, media, and self-reasoning; (ii) realized performance evolving with learning and scale; and (iii) adoption as a purchase rule.

We showed how a kinetic model delivers a tractable two-state system for adoption, performance, and perceived value. In the deterministic limit, the purchase rule reduces to a hard threshold, yielding piecewise exponential adoption and sharp, interpretable conditions for a hype peak and a trough: expectations fall when media/content unwind faster than performance. Re-introducing heterogeneity via the Fokker-Planck stationary distribution yields a smoothed adoption hazard—recovering the empirically observed S-curve—while preserving the same mechanism and inequality structure for the onset of disappointment. We further showed, via a concise DSMC agent-based simulation, that the mechanism is robust to relaxing time-scale separation and even persists under constant media: endogenous feedback and heterogeneity suffice for HC emergence. Taken together, these results provide (i) a micro-to-macro derivation of the hype cycle, (ii) explicit, testable conditions for its appearance, and (iii) a coherent way to interpret policy or managerial levers. The model yields an explicit condition for the emergence of a hype peak, i.e. equation 4.20: market expectations rise while media influence grows faster than adoption and peak once adoption growth and performance catch-up dominate, producing a correction when promised value exceeds realized performance. This condition is empirically testable using proxies for media exposure and message content, adoption rates, and objective performance indicators. The same structure provides an inter-

pretable set of levers: increasing media conveyed value w can accelerate early takeoff but also increases overshoot risk when $w - r$ is large; lowering the effective price P (via pricing or subsidies) shifts the adoption threshold; and increasing the responsiveness of investment to adoptions (i.e. k_1) mitigates disillusionment by narrowing the expectation-performance gap. These mappings suggest concrete managerial and policy trade-offs between accelerating diffusion and avoiding excessive expectation overshoot.

The approach has limits. First, the realized performance function $r(t)$ is specified in a modular way (sigmoid in cumulative perceived value and adoption); while convenient for analysis, the exact form is not pinned down and is likely technology-specific. More empirical and theoretical work is needed to identify functional classes for $r(\cdot)$ that match engineering evidence (S-curves with shifting ceilings, scale-learning terms, or path dependence) and to determine when different forms qualitatively change the dynamics. Second, in the current presentation we normalize or exogenize some drivers (e.g., media content level), which helps isolate the core mechanism but understates strategic behavior and endogenous mechanics in media behavior. Third, identifiability is nontrivial: media, adoption, and performance are intertwined, so parameter inference will require careful priors and multiple data sources.

In sum, the chapter advances the literature by turning the HC from a descriptive curve into a derivable, analyzable phenomenon with clear microfoundations and falsifiable conditions. Its strength is the combination of analytical clarity and modularity; its main weakness is the current agnosticism about the precise form of $r(t)$. Resolving the latter will turn this from a compelling theory into a practical toolkit for anticipating and managing hype–disappointment dynamics.

The DSMC scheme employed in this chapter to numerically solve the system provided a very simple example of how agent-based simulations can be used to explore with great flexibility regimes that are not easily fathomable with analytical tools. In the next chapter, we present another example of ABM applied to innovation diffusion. Using the data from the SINFONICA European Project, we show how ABMs can reflect the complexities of reality and model heterogeneity at the maximum granularity.

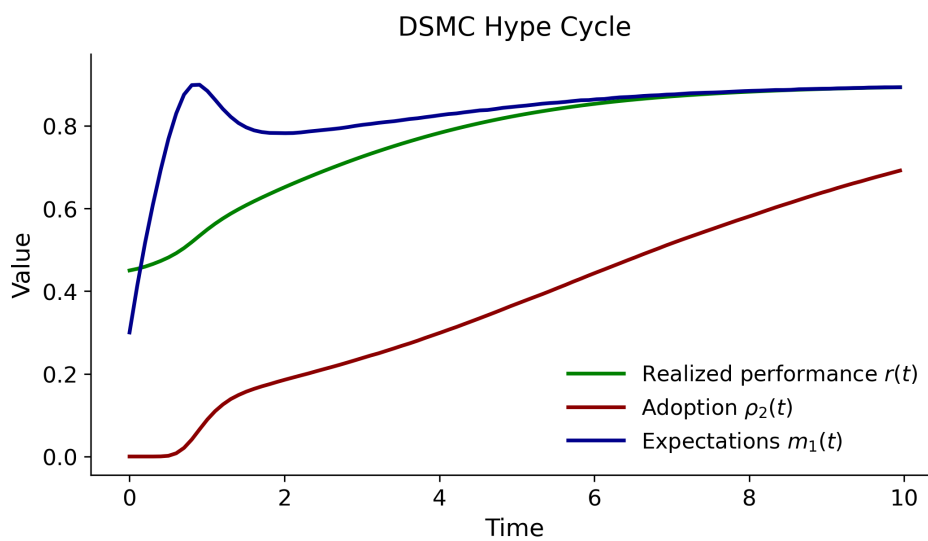
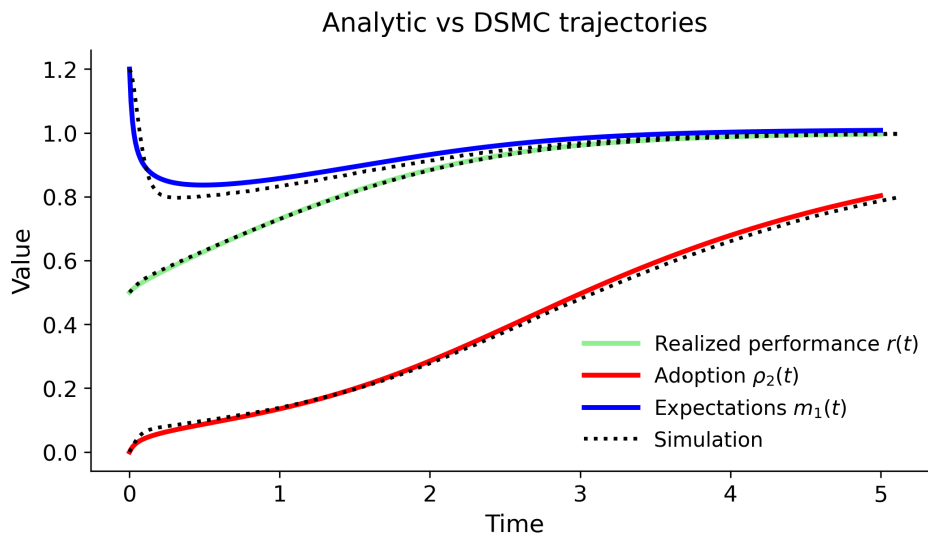


Fig. 4.3 The evolution in time of expectations (blue), adoption (red) and performance (green) obtained with the DSMC.

Chapter 5

An Agent-based Model for the Diffusion of Automated Vehicles

Having moved from descriptive S-curves to micro-founded dynamics, we now lean fully into computation. This chapter, following Masali and Perboli (2024), sets aside analytical tractability to leverage modern computation and rich data, building an agent-based model that mirrors the heterogeneity of real populations – who people are, what they value, and how they respond to experience and information. The goal is not a new closed form, but a faithful playground: a place to watch diffusion play out under plausible rules and to see how small differences in people and circumstances add up to the large, uneven patterns we measure.

5.1 General Context

The concept of a smart city plays a central role in addressing urban challenges and complexities through digital transformation strategies. This necessitates active involvement from all stakeholders, ranging from the private and public sectors to city residents, alongside municipal bodies (Perboli and Rosano, 2020).

The rapid evolution of urban environments into smart cities is driven by the need to improve efficiency, sustainability, and quality of life for their inhabitants. Central to this transformation is the integration of innovative transportation services that leverage cutting-edge technologies (Apruzzese et al., 2023; Brotcorne

et al., 2019). One such innovation is Connected, Communicating, and Automated Mobility (CCAM) systems, which promise to revolutionize urban transportation by enhancing connectivity, reducing congestion, improving safety, and lowering pollutant emissions. CCAM systems rely heavily on advanced technologies such as edge and fog computing, Artificial Intelligence, and IoT to process vast amounts of data in real-time, enabling vehicles to communicate with each other and with urban infrastructures. These technologies collectively support the complex data flows and real-time decision-making required for effective CCAM deployment.

In this chapter, we present an agent-based model (ABM) that simulates the diffusion of CCAM in the context of the Noord-Brabant region of the Netherlands. The model is informed by data from the SINFONICA project, the flagship European initiative aimed at facilitating the deployment of innovative CCAM services in an inclusive and equitable way (SINFONICA Consortium, 2023). By characterizing agents based on sociological, demographic, and behavioral data, our model provides insights into how CCAM can be adopted and scaled in urban and suburban environments. The findings from these simulations offer valuable implications for city planners, policymakers, and technology developers working towards the realization of smart cities.

Agent-based modeling has been extensively employed in the study of both transportation and innovation diffusion, as it allows to model the behavior of people in the population at the individual level. Here, we propose a model that merges these two aspects and captures both the transportation-related behavior and the consumer behavior of the population. Notable works that employed ABM to predict the diffusion of innovative products include the application to the diffusion of photovoltaic solutions in Italy (Palmer et al., 2015; Schiera et al., 2019); diffusion of energy-saving behaviors (Bastani et al., 2016; Hesselink and Chappin, 2019); and diffusion of electric cars (Kangur et al., 2017). The authors in Shi et al. (2023) employed ABM to create a framework to predict the diffusion of green technologies in general, with a particular focus on the agents' sentiment towards innovation. In this work, the sentiment towards CCAM will play a pivotal role, and the evolution of that sentiment recalls the framework of the the Hype Cycle (Dedehayir and Steinert, 2016), also subject of Shi et al. (2023).

The ability to model traveling habits at the individual level has made ABM popular in the transportation community as well. Many studies have focused on

modeling various aspects of automated vehicles on the economic and urban system (Boesch et al., 2016; Fagnant and Kockelman, 2014; Lokhandwala and Cai, 2018). For a comprehensive review of the literature in this field, see Bastarianto et al. (2023).

The chapter is organized as follows: Section 5.2 explains the methodology employed for both the data analysis and the creation of the agent-based model. Section 5.3 presents the results obtained for some specific configurations of the model parameters. Finally, Section 5.4 concludes the chapter.

5.2 Methodology

5.2.1 Data

The data and the problem setting used come from the context of the European-funded project SINFONICA (SINFONICA Consortium, 2023). The data collection consisted of interviews with people from four different cities or provinces in Europe (Trikala, Greece; Hamburg, Germany; West Midlands metropolitan area, UK; and Noord-Brabant province, Netherlands). The model is built using only the data from the Noord-Brabant province, for consistency. The interviews were conducted on small focus groups with diverse socio-economical background, and the subjects could express personal comments on multiple aspects of the questionnaire. The questionnaire is divided into four main sections:

- Socio-economical information, such as age, gender, education level, and income;
- Current transportation habits, priorities and limitations. The possible priorities were time savings, availability at any time, safety and security, travel cost, reliability, environment, reachability, comfort, health-related issues and cleanness;
- Inclination toward innovation and technology and digital skills;
- Opinion on CCAM.

Many questions allowed for a 5-point Likert scale response. Firstly, preliminary data cleaning and preparation have been conducted on the dataset. We discarded

all personal comments and extra information that were too specific to be used to characterize agents in an agent-based model. We then re-arranged information on preferred modes of transportation to isolate the two overall favorite modes for each interviewed person. The final dataset consisted of 39 items with a list of 42 relevant attributes. In order to build an informative agent-based model, we enlarged the sample with a synthetic population generation procedure based on statistical analysis and clustering. In particular, we standardized the dataset and applied a Gaussian Mixture clustering algorithm (Xu and Wunsch, 2005) to the standardized data, with the creation of tuples of attributes to preserve the statistical correlation among key columns, such as the ones containing information on the preferred modes of transport and the travel priorities of people. To optimize the number of clusters, we used the Bayesian information criterion (BIC) (Neath and Cavanaugh, 2012). The optimal number of clusters is the one producing the lower BIC score. Figure 5.1 shows that the optimal number of clusters in the dataset is four. To obtain the

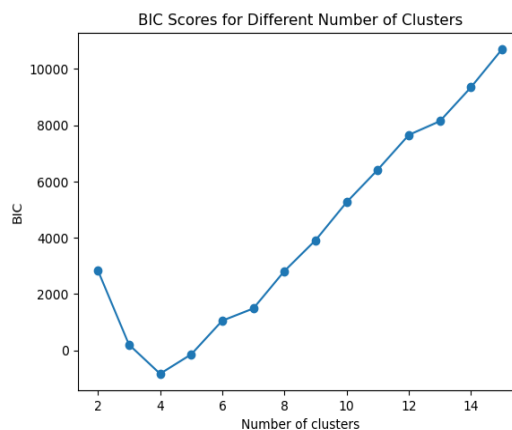


Fig. 5.1 BIC score for different numbers of clusters.

synthetic population, we generated new instances by sampling from the multivariate normal distribution fitted to the numerical data within each cluster, or based on the value distribution for categorical attributes. Finally, we deleted outliers and unrealistic instances. The synthetic population we ended up with is composed of 10000 individuals with a set of attributes statistically distributed as the attributes in the original sample. We used this population to characterize the agents of an agent-based simulation of the adoption of CCAM.

5.2.2 Agent-based model

To build an agent-based model from the data collected through the SINFONICA survey, we needed to define behavioural rules based on the available information. As a preliminary informative tool, we ran some statistical tests on the attributes to assess correlations among them. In particular, we created a correlation matrix among numerical values, that did not reveal strong correlations. For categorical values we performed a chi-square test that unsurprisingly showed correlations among attributes that refer to the same domain. In our model, agents represent individual city residents who interact with their environment and make decisions based on predefined rules. Each agent is characterized by socio-economic attributes, behavioral attributes, technological attributes and sentiment attributes. The decision-making process for adopting CCAM is influenced by factors such as perceived benefits, social influence, innovativeness and accessibility. Figure 5.2 shows the attributes of the synthetic

ID	Age	Gender	# Cars	# Mbikes	# Bikes	Education	Employment	Ppl in household	Children	Income	Nationality	Area		
ID	Age	Gender				Education	Employment	Ppl in household	Children	Income	Nationality	Area		
Mode 1	Mode 2	Distance	Priority 1	Priority 2	Priority 3	Satisfaction 1	Satisfaction 2	Limit1	Limit2	Limit3	Limit4	Limit5	Limit6	Limit7
Mode 1	Mode 2	Distance	Priority 1	Priority 2	Priority 3	Satisfaction 1	Satisfaction 2						Health issues	Digital limit
PC	SmartPh	App usage	Tech trust	Tech interest1	Tech interest2	Tech interest3	Tech interest4	Tech interest5	Tech confidence	Data attitude				
Access	Tech attitude													
CCAM awareness	CCAM past use	CCAM sentiment												
CCAM awareness		CCAM sentiment	m1p1	m1p2	m1p3	m2p1	m2p2	m2p3						

Fig. 5.2 The attributes of the survey data (first and third lines) and the corresponding attributes of the agents (second and fourth lines).

population and the attributes of the agents that were calculated from them, different colours correspond to different domains. The socio-economic data (in blue in the figure) were directly assigned to the agents, except the number of vehicles owned, which has been neglected. The behavioral data (in red in the figure) were also transferred unchanged, except data on the limitation to mobility, that are highly correlated with the priorities and so have been neglected. The last three questions regarding the limitations concerned health problems and digital limitations and have been included into the “Health issues” and the “Digital limit” attributes. The technological data (in violet in the figure) have been condensed into the two attributes

Priority	Walking	Bike	Car	Publ. transp.
Time savings	5	20	30	15
Availability at any time	0.5	0.5	1	0
Safety & Security	1	0	0.5	0.5
Cost	0	0	150	100
Reliability	1	0.5	1	0.5
Environmental reasons	1	1	0	0.5
Reachability	1	1	1	0.5
Comfort	0.5	0.5	1	0.5
Health-issues	0	0	1	0.5
Physical activity	1	1	0	0.5
Cleanness	1	1	1	0.5

Table 5.1 Satisfaction thresholds for mode-priority pairs.

“Access”, a binary attribute that encodes if the agent has access to any type of digital device, and “Tech Attitude”, which contains information on the attitude of the agent towards innovation and new technologies and is a number between 0 and 1. The awareness of the existence of CCAM and the sentiment towards it (in yellow in the figure) are contained into the “CCAM awareness” binary attribute and the “CCAM sentiment” attribute, respectively. Lastly, six attributes (in green in the figure) are calculated that encode the performance threshold that the innovative CCAM vehicles have to reach in order to satisfy the agent. These attributes depend on the preferred transport mode and the travel priorities of each agent, plus information on her or his health issues. In this work, they have been calculated following table 5.1.

These are key attributes that determine the shape of the diffusion curve. If data on the actual diffusion of CCAM will be available, these attributes calibrated in order to reproduce the actual diffusion curve could give insights on the actual performance requirements of the population. In this study, the parameters have been assigned placeholder values, with the primary focus being on the relationships between them. Note that “Health issues” and “Physical activity” both correspond to the “Health-related issues” option for the priorities: in one case the “Health issues” attribute of the agent is 1, in the other case it is 0. Hence, each agent is equipped with three threshold values corresponding to one of its preferred modes and three corresponding to the other, all calculated according to its travel priorities. The CCAM service will be characterized by a score in each priority category and this will determine its actual diffusion in the area. CCAM service will also have a minimum digital skill

requirement to be used. An additional attribute of the agents that does not come directly from the survey data is the status of the agent, that gives information on the fact that the person has not adopted yet, has adopted but is not convinced yet, is a convinced user or has abandoned the innovation. The actual agent-based model has been developed using the software NetLogo (Wilensky, 2015). This software allows to simulate the behavior of an adequate number of agents and to easily derive insights on the evolution of the system by plotting graphs and computing statistics. If the number of agents gets significantly higher than 10000, it might be too slow to run a simulation on NetLogo. Once the world is populated with 10000 agents with the described attributes, interaction and behavioral rules have to be implemented. We created four different mechanisms that drive the behavior of the agents and the diffusion of the innovative transport system, namely interaction, advertisement, purchase, and abandonment. At every simulation time-step (or tick), all the agents that are *non-users* interact with another agent and possibly purchase, all the agents that are *on trial* decide whether to become *users* or abandon the innovation and become *haters*. The interaction mechanism controls how agents interact with each other and how this interaction changes their beliefs and status. In this work, a social network structure that underlies the interactions between agents is not present, as no information about social structure was present in the SINFONICA dataset: each agent can equally interact with every other agent in the simulation. It would be interesting to enrich the model with a network structure that could create bubbles of faster or slower diffusion or perhaps reveal key nodes with higher influence. When an interaction between agent A and agent B occurs, agent A can change some of its attributes:

- If agent A is not aware of the innovation and agent B is, then agent A becomes aware.
- If agent A has a sentiment towards CCAM lower than the maximum value (which is 4) and agent B is a *user*, then agent A increases its sentiment towards CCAM by 1.
- If agent A has a sentiment towards CCAM higher than the minimum value (which is 0) and agent B is a *hater*, then agent A lowers its sentiment towards CCAM of 1.

- If agent A has access to a digital device and it is aware of the existence of CCAM and if agent B is a *user*, then a purchase event occurs.

The fact that the purchase may occur only after the interaction with a *user* models the “word of mouth” effect. The purchase mechanism controls how *non-users* go into *on trial* period. If the digital skill of the agent is higher than the required digital skill the agent might adopt the innovation: this will depend on its level of satisfaction (0, 1 or 2) with one of its current transport modes and on its sentiment towards CCAM (0, 1, 2, 3 or 4). These attributes have been built so that if the sentiment is higher than the satisfaction level with one of the modes +1, the agent replaces that mode with CCAM with a probability that is proportional to its “Tech attitude”. If this happens it also changes its status from *non-user* to *on trial*. This mechanism has been thought so that an agent that is not satisfied with its current transportation habits is more likely to try the new system even if he or she has some reservations on CCAM. Furthermore, the adoption of the new technology is more likely the higher the agent’s inclination towards technology is. The advertisement mechanism controls how agents interact with marketing campaigns and how higher investments in advertisement can influence the diffusion of an innovation. Every time an agent interacts with advertisement, the same thing as if it interacted with a *user* happens and a purchase event might occur. The probability of an agent interacting with advertisement is proportional to parameter ADV-strength, which represent the investments and effectiveness of the advertisement campaign. The last rule that governs the agents’ behaviour in the model is the abandonment rule. After a trial period, newcomers decide whether they like or not the CCAM system and become either *users* or *haters*. This will happen based on the performance levels of the system and on their personal performance thresholds. Given a transportation mode that has been replaced with the CCAM system, each agent has three threshold values associated with it. The algorithm computes the number (0, 1, 2 or 3) of requirements that the CCAM system fulfills, according to its performance attributes. The agent will decide whether to buy or not based on its satisfaction level with the old transportation mode and on the number of requirements that the new system satisfies. In particular:

- If the CCAM service satisfies all three requirements the agent becomes a *user* no matter the satisfaction level with the old mode.

- If the CCAM service satisfies two requirements the agent becomes a *user* if the satisfaction level with the old mode is 0 or 1. It becomes a *hater* otherwise.
- If the CCAM service satisfies only one requirement the agent becomes a *user* if the satisfaction level with the old mode is 0. It becomes a *hater* otherwise.
- If the CCAM service does not satisfy any requirements the agent becomes a *hater* no matter the satisfaction level of the old mode.

These four rules enable the agents that populate the simulation to interact with each other and with the environment, to create their own ideas on the new transportation system and to act accordingly, adopting or not the innovation based on their priorities and the performance level of the technology. The simulation allows to reproduce various scenarios of technology readiness and advertisement investment and to compute various performance indicators, such as the number of users of the CCAM service, the number of users of each traditional transportation mode, the average sentiment towards CCAM. In the next section, we will show the outcomes of some of these scenarios and discuss the potentiality of the model for policy makers and stakeholders.

5.3 Results

The scope of this chapter is to present the agent-based model that we have built based on the dataset of the SINFONICA project, and to show some of its potentialities. The results we describe in this section are aimed to illustrate different kinds of scenario that can be tested inside the agent-based framework. Let us stress here that the threshold levels we set following table 5.1 are important parameters that control the outcome of the simulation. In this case, they have been set to placeholder values and the only important thing is where, between the classical transportation modes, the innovative system positions its performance parameters. We will illustrate two different scenarios, with different combinations of parameters, that will show qualitatively different diffusion processes.

5.3.1 Scenario 1

With the first scenario, we are going to investigate the effects of different average speeds of the CCAM system on its diffusion. In the dataset we are using, “time savings” is the most common priority indicated by the respondents, both overall and by car users (which the SINFONICA project aims to reduce). For this purpose, let us set the parameters of the technology to mimic the performance level of an excellent public transportation service. This means to set the parameters according to the first column of table 5.2 and to vary the average velocity to outperform traditional transportation modes, progressively. CCAM speed values are set to 10, 18, 25 and 35, respectively.

Priority	Scenario 1	Scenario 2
Speed	-	10
Availability at any time	0.8	0.4
Safety & Security	0.8	0.4
Cost	90	160
Reliability	0.8	0.4
Green-nes	0.8	0.4
Reachability	0.8	0.4
Comfort	0.8	0.4
Physical activity	0.8	0.8
Cleanness	0.8	0.4
ADV strength	0.0015	0.004

Table 5.2 CCAM system parameters for different simulations.

We show in figure 5.3 the diffusion curves of this scenario, along with the evolution of the number of car users. This is an example of progressively more successful diffusion processes. In the final case the new technology is accepted by the population, there is a good number of *users*, a contained number of *haters* and every person has tried the innovation before deciding whether to adopt it. The small number of non users that had not tried it are persons without access to a digital device, that cannot adopt the innovation. If we observe the evolution of the number of users of the various modes, it is easy to see how the car users are the category that is the most disappointed in the CCAM service and how the car levels return almost to initial levels after a short drop in users that were trying the new service. In fact, the percentage of old car users among the adopters compared with the initial number is the lowest, together with the number of walker adopters (who make up

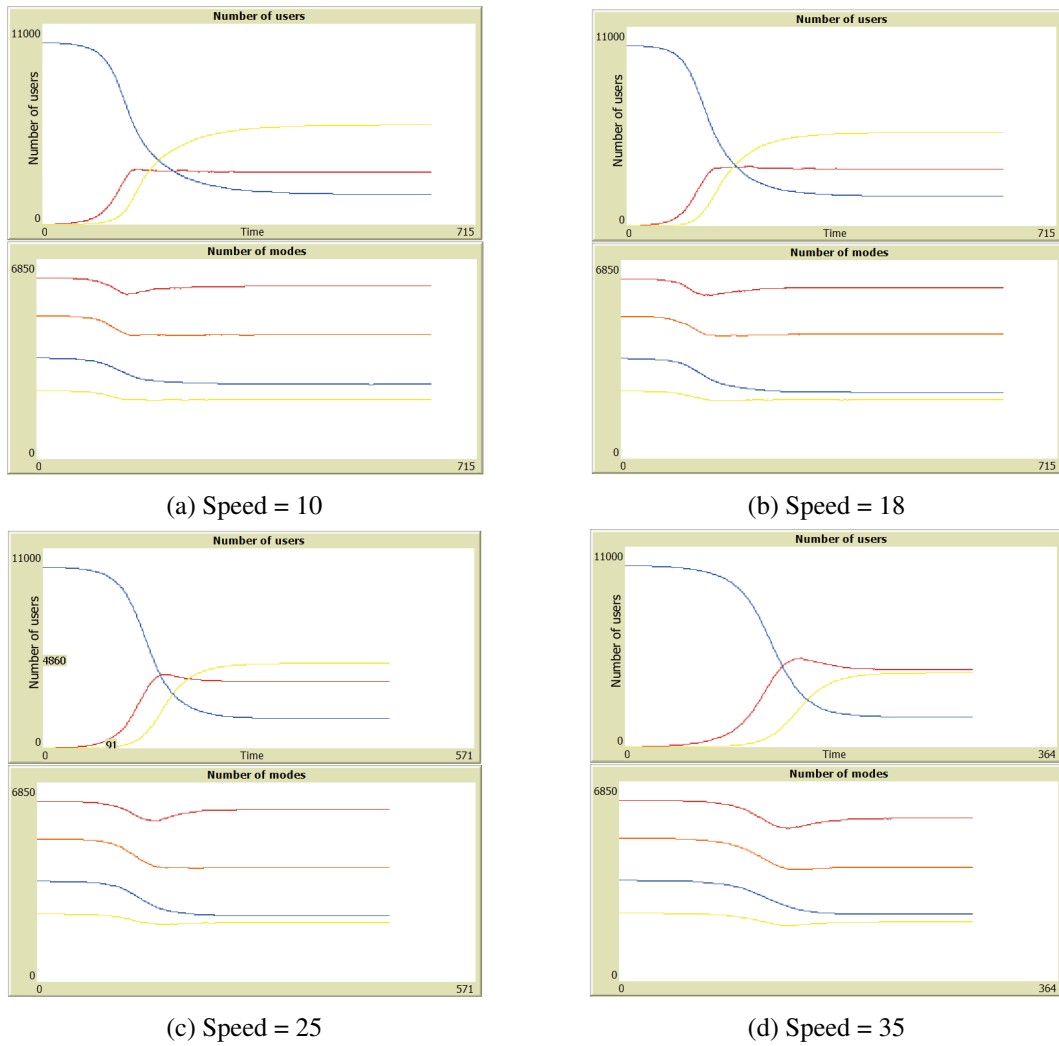


Fig. 5.3 Evolution of the diffusion process with number of users (*non users* in blue, *users* in red and *haters* in yellow) and travel mode usage (car users in red, bike users in orange, public transport users in blue and walkers in yellow), for different values of Speed of the CCAM system.

the smallest fraction of users and are of little concern). This indicates that, according to the dataset from SINFONICA and to this model, reducing car levels in the area is a difficult task and that car users are not easily convinced to switch to a more sustainable mobility. To design a CCAM public transport system that appeals to car users would mean to increase the reachability and the comfort of the service, which are the most common priorities of car users according to the interviews, as well as its velocity. This will bring up the levels of old car users among the adopters and consequently reduce the number of cars in the area. The agent-based model presented in this chapter enables decision makers to test different configurations of parameters to understand the effects on the diffusion process of different product designs, or communication campaigns.

5.3.2 Scenario 2

The Bass model (Bass, 1969) divides the adopters of an innovative product into innovators and followers, where innovators independently adopt the innovation, while followers adopt under the pressure of a “social force”, or “word of mouth”, that grows as the number of users grows. Arguably, the innovators do not act autonomously but are convinced to purchase by the communication campaign that the company selling the innovation has funded. The second scenario aims to investigate the effects of advertisement on the shape of the diffusion curve and on the evolution of the mean sentiment towards CCAM in the population. As the effectiveness or the magnitude of the communication campaign grows, the less important the social force effect becomes, and the S-shaped diffusion curve straightens. Furthermore, if the technology is not ready, advertisement may drive the expectations towards a technology up to a point where they are not met by the actual performance of the product, and disappointment might happen, letting the Hype Cycle emerge.

Figure 5.4 shows the evolution of the average sentiment in the population towards CCAM and has been obtained by setting the parameters as in the second column of table 5.2, corresponding to a bad service. The strength of the advertisement campaign has been set to 0.004, higher than the previous case. The oscillating shape of the curve shows that the agent-based model is able to recreate, under specific circumstances, the hype phenomenon. Note that, if the performance level of the new system is not high enough, a subset of agents does not adopt it in the long run without having tried it, but because they were told of the poor performance and have a low




Fig. 5.4 The evolution of the average sentiment following the Hype Cycle pattern.

sentiment towards the technology. This is observable in the model even without an underlying social network structure, that may enhance this effect. For example, the simulation that produces the hype cycle effect described above ended with 3104 *non-users*, of which 1462 with low sentiment value and access to digital devices.

5.3.3 Scenario Generation

The scenarios above are illustrative, by design. Our only empirical foundation is the population characterization from SINFONICA's Noord-Brabant focus groups: a small but carefully cleaned sample that we expanded into a 10,000-agent synthetic population preserving the joint distribution of key attributes. What we do not observe are two crucial objects: (i) the realized performance trajectory of a CCAM service that would actually operate in the region, and (ii) an observed adoption path for CCAM to calibrate the diffusion curve. In the absence of those data, the model is used as a controlled thought experiment: we hold the empirically grounded population fixed and vary a small set of CCAM service attributes and communication intensity to examine how a plausible, data-informed society would react. The thresholds in Table 5.1 and the service parameters in Table 5.2 are therefore placeholders chosen for internal consistency and contrast, not measurements of any existing deployment.

This choice is consistent with the purpose of this chapter. The goal here is not to forecast the exact adoption of a specific CCAM product, but to map mechanisms: how sentiment, word-of-mouth, and trial-abandon decisions interact with heterogeneity in priorities and access; how raising one performance dimension (e.g., speed) shifts adoption across user types; how stronger advertising can straighten an S-curve and amplify disappointment when performance lags. The scenarios are counterfactual probes that show what the empirically characterized population would do under

specified, hypothetical technology profiles. When richer data on CCAM performance and early adoption become available, the same framework can be re-estimated, the placeholder parameters replaced with measured ones, and the exercise upgraded from illustrative exploration to calibrated policy analysis.

5.4 Conclusions

In this chapter, we have given an example of the potentiality of an agent-based model for the diffusion of an innovative transportation system. The model has been populated with agents characterized by attributes that define both their travel and their consumer behavior. The attributes are informed by data from the SINFONICA project, a European funded project with the goal of facilitating the adoption of Connected, Communicative and Automated Mobility (CCAM) systems in Europe. The SINFONICA dataset consisted in interviews with a panel of 39 people from the Dutch region of Noord-Brabant. The dataset has been cleaned and analyzed and a 10000 people synthetic population has been generated, maintaining the same statistical properties of the original sample. The agents in the model have attributes defining their socio-economic status, their preferred travel mode, their travel priorities, their technological skill level, their inclination towards innovation and their sentiment towards CCAM. Once the simulation is populated by the agents, we set the rules that control the different behaviour of those agents, namely interaction, purchase, advertisement and abandonment. The agents interact between one another and exchange information and opinions on the innovative system and they change their sentiment towards it. The advertisement campaign is another mechanism that can influence the opinion of the population on the product. The agents then decide whether to try the innovation based on their level of satisfaction with their usual modes of transportation and their sentiment towards the new mode. The purchase is always triggered by the interaction with the advertising campaign or with a fellow agent who is a user of the innovation. This ensures the model to account for both spontaneous innovation and word of mouth effect. After a period of trial, the agents decide whether to abandon the innovation and return to their old mode of transportation or to integrate the new mode into their habits. This decision is taken based on the performance level of the service and the satisfaction of the agent with its usual modes. If an agent decides to abandon, it becomes an *hater* and spread bad

reviews of the technology in the population. We showed two possible scenarios for the diffusion of CCAM in the area. The first scenario is an example of successful diffusion, where the technology is well accepted by the population. We show how the car users are hardly seduced by the CCAM service and that the ABM can provide insights for informed product design and/or communication strategies. The second scenario shows the effects of advertisement and poor performance product on the diffusion process. We showed how, under specific circumstance, the evolution of the mean sentiment towards the innovation in the population follows the Gartner hype cycle, rising to a peak of inflated expectations and then falling when compared to the actual performance of the service. This aspect of the model has not been thoroughly investigated and is an interesting future development worth pursuing. In general, the model is able to reproduce a variety of qualitative diverse adoption patterns and can be a precious tool for decision makers to test different diffusion scenarios. Future improvements include the introduction in the model of the social network structure. In addition to that, the dataset it was built on is limited. In the future, SINFONICA will provide larger datasets and we will be able to better characterize our agents and to run more sophisticated simulations.

Chapter 6

Conclusions

In this dissertation, we have investigated how the lens of complexity can improve our understanding of how innovations diffuse in societies. This approach places the focus on the hidden mechanics that govern the observable outcomes, on the interactions, the feedbacks, the non-linearities that shape the spread of new things, and their acceptance. In practical terms, that means treating aggregate S-curves as summaries rather than mechanisms, and making explicit the micro- and meso-processes that bend those curves: heterogeneous agents exchanging information with unequal credibility, media signals rising and waning, and realized performance catching up (or failing to) with early claims. When these ingredients are modeled explicitly, familiar deviations – stalls, overshoot, and “hype-disillusionment” – are not anomalies but expected regimes of a multi-scale system. The pay-off is twofold: cleaner explanations and decision-relevant forecasts that respect uncertainty and identification limits. Accordingly, the thesis progressively zooms in on the underlying mechanics of diffusion. In the following, the main contributions from each chapter are clearly stated.

6.1 Main Contributions

In chapter 2, we showed that a disciplined adaptation of the Bass model can summarize national progress toward the SDGs and extract timing diagnostics (take-off/peak improvement/saturation) from country trajectories. The contribution is intentionally modest and practical: treat the S-curve as a summary of a many-actor system, not as

a mechanism. The analysis highlights heterogeneity across countries—consistent with structural capacities and constraints—and argues that, despite shocks and policy complexity, a compact curve still travels well enough to support comparable timing forecasts.

In chapter 3, we moved from description to mechanism by deriving an operative adoption equation from a kinetic (Boltzmann) model of interacting, boundedly rational agents. The closed-form law (with the regularized lower incomplete gamma) links macro S-shapes to micro updating rules and keeps interpretable levers – media strength w and realized performance r – explicit. Empirically, the model matches classic benchmarks (TV, AC, dryer, freezer) and is shown to be suitable for early-stage forecasting: using only the first 3–5 years of data, it delivers planning quantities (market potential, peak timing/height, saturation) with quantified dispersion; results improve with five points and are product-specific where uncertainty remains. Distributional fitting via Bayesian inference of the parameters is also performed to provide a more informative and reliable fitting of the empirical data. The kinetic innovation diffusion model is framed as a bridge between aggregate Bass-type models and simulation-only ABMs.

In chapter 4, we extended the lens from adoption to expectations, coupling perceived value, adoption, and performance to turn the Hype Cycle from a picture into a set of testable conditions. The chapter states the peak condition linking media and adoption growth to performance catch-up, shows how heterogeneity smooths adoption via the kinetic/Fokker-Planck route, and uses a DSMC scheme to confirm that hype-like dynamics arise even with fixed media once the feedbacks are in place.

In chapter 5, we provided a data-informed ABM of CCAM diffusion to show what a micro-level, policy-tunable model can deliver. Using SINFONICA interviews from Noord-Brabant, we built a 10000-agent synthetic population via Gaussian mixtures, we endowed agents with socio-economic and travel attributes, technology attitude, and sentiment, and we implemented interaction, advertising, purchase, and abandonment rules in NetLogo. Two scenarios illustrate qualitatively distinct regimes. With strong performance, diffusion succeeds and sentiment stabilizes; with weaker performance and stronger advertising, the mean sentiment oscillates in a pattern consistent with hype and disillusionment; car users prove hardest to convert, invoking ad hoc strategies to ensure a successful diffusion process.

6.2 Limitations

The methodologies presented in the dissertation have limits that must be taken into consideration and are tied to the scale and purpose for which they are intended. The value of the chapter 2 is descriptive: it treats the S-curve as a compact summary of multi-actor dynamics, not a structural mechanism. That choice has trade-offs. First, results depend on preprocessing choices (shifts of the initial condition, smoothing in noisy series, Monte-Carlo restarts) and on guardrails like forcing $M = 100$ in near-linear cases; these steps are necessary for stable timing diagnostics but also introduce operator judgment. Second, the framework cannot accommodate structural breaks or regime changes except by excluding non-S trajectories or refitting; as a result, countries with instability or shocks are precisely those where the method is least informative. Third, because the Bass layer is mechanism-agnostic, it cannot by itself attribute changes in slope to policy, institutions, or network effects.

Chapter 3 explains the S-shape from micro rules and yields a closed form, but it pays a cost in identifiability. Early in diffusion, different parameter bundles (e.g., media strength w vs. interaction parameters, or total number of agents N vs. steepness via γ) can fit the same partial trajectory; point estimates risk overconfidence. Interpreting fitted parameters as literal behavioral quantities can be dangerous and ensembles are preferred over single fits. A second limitation is model parsimony in heterogeneity of the agents: multi-class extensions (different a_i, b by segment) are natural but can compromise the neat closed-form hazard that makes equation 3.46 tractable, and have not been implemented in the dissertation. Finally, mapping real advertising effort into w is out of scope, so w remains an abstract control rather than an operational spend-to-impact function.

Chapter 4 isolates a clear mechanism and delivers testable peak-trough conditions, but it does so by making cleaning assumptions. Most importantly, the realized-performance path $r(t)$ is kept generic and its form is not thoroughly investigated; the analytic clarity this buys comes at the price of technology-specific fit and leaves open which functional classes of $r(\cdot)$ best align with engineering evidence. In places, media intensity is normalized or treated as exogenous to focus on the mechanism; this tempers claims about strategic media behavior and co-movement with adoption. Finally, joint identification of media, adoption, and performance is hard: without informative priors and multiple data sources, parameters will be entangled.

The model described in chapter 5 is informative as a calibrated playground, but several limits remain. The thresholds that gate adoption/abandonment are placeholders in the illustrative runs; different settings would change quantitative outcomes. Interactions are well-mixed in the reported experiments (no explicit network), so cascade patterns and targeting strategies are only partially explored. External validity is inherently bounded by the data: the synthetic population is learned from a small, local sample (Noord-Brabant focus groups), and larger SINFONICA waves are needed to refine agent heterogeneity and to test portability. Finally, the mapping from advertising intensity to outcomes is stylized (a single “ADV-strength” knob), which is adequate for scenarios but not yet a market-ready media model.

6.3 Future Developments and Final Remarks

The results obtained in this dissertation suggest several natural continuations. They all follow the same logic that guided the thesis: use each layer for what it does best, and each tool for the right purpose.

A first direction is to extend the investigation on the relation between NIEs and SDG development by leveraging tools such as agent-based and network diffusion models. This would allow to model explicitly the interactions – funding ties, knowledge flows, supply-chain links, public attention – that happen among the heterogeneous actors of innovation ecosystems – ministries, agencies, regions, firms, universities, financiers, NGOs, media, citizen groups. This would mean to build a multilayer network over which capabilities, practices, and technologies diffuse. Agents would follow simple decision rules – investment, collaboration, adoption, imitation, lobbying – and the synergies, bottlenecks, and feedback loops would emerge from how layers interact. This would allow to ask structural questions that a Bass fit cannot answer: which missing connections produce stagnation, which coalitions unlock SDG progress, and where coordination failures create fragility in otherwise promising trajectories.

A second line concerns the extension of the kinetic diffusion model to move it closer to real markets. To incorporate richer heterogeneity of agents, design more realistic engaging mechanisms with the innovation (abandonment, repeated purchases), or adopting an endogenous media signal would go in that direction. Heterogeneity can be encoded in a finite set of segments – enthusiasts, pragmatists, laggards – with

segment-specific credibility, media sensitivity, and thresholds. The macroscopic law then becomes a mixture of segment hazards that still retain the regularized-gamma structure. Engagement should also move beyond one-shot adoption: add abandonment so the adoption stock can contract when expectations fall or performance disappoints, and introduce an age-dependent replacement process to separate penetration from sales and produce realistic multi-peak sales without inflating ultimate adoption. A light multi-unit layer covers categories with repeated purchases. Finally, one could close the media loop by making attention an endogenous state that responds to adoption acceleration and performance improvements.

A third direction is to create a data-informed realized performance function $r(\cdot)$ to plug into the system of chapter 4 and to collect coherent data to inform a useful fitting of that model. This would turn $r(\cdot)$ from a placeholder into a data-informed driver, tightening identification and making the Hype Cycle mechanism estimable and decision-relevant.

Finally, an immediate extension to the ABM described in chapter 5 is to embed agents in explicit networks (small-world, modular, spatial) and to study how topology changes cascade patterns, targeting efficiency, and the conditions under which hype-like sentiment persists or is damped. The synthetic population pipeline is already in place; adding network structure and running sensitivity analyses would test robustness without altering the behavioral rules.

This work closes where it began: with the claim that diffusion belongs to the natural habitat of complex systems. Heterogeneous actors, interactions, and nonlinear feedbacks generate patterns that no top-down curve can fully explain on its own. The thesis responded by mixing parsimony and mechanism – keeping summaries where they are informative, and bringing in structure where they are not. In doing so, it aligned innovation studies with a broader tradition in complexity science that moves smoothly between microscopic rules and macroscopic emergent behavior, without losing rigor and predictive discipline.

The result is not a unified theory but a useful set of instruments: innovation diffusion should be investigated with the tools of complexity, matched to the scale of the question and tied back to evidence. National trajectories can be summarized; expectations and performance can be modeled as coupled dynamics; ecosystems can be explored as interacting heterogeneous agents. Together, these choices make

diffusion intelligible and actionable, acknowledging limits while extracting the structure that matters for policy and management.

References

- Abdullai, L., Porras, J., and Haque, S. (2022). Building a national smart campus to support sustainable business development: An ecosystem approach. *arXiv preprint arXiv:2209.13613*.
- Agasty, S., Tarannum, F., and Narula, S. A. (2023). Sustainability innovation index for micro, small and medium enterprises and their support ecosystems based on an empirical study in india. *Journal of Cleaner Production*, 415:137793.
- Albitar, K., Al-Shaer, H., and Liu, Y. S. (2023). Corporate commitment to climate change: The effect of eco-innovation and climate governance. *Research Policy*, 52(2):104697.
- Alkemade, F. and Suurs, R. A. (2012). Patterns of expectations for emerging sustainable technologies. *Technological Forecasting and Social Change*, 79(3):448–456.
- Ancona, A., Cinelli, M., Ferraro, G., and Iovanella, A. (2023). Network-based principles of entrepreneurial ecosystems: a case study of a start-up network. *Small Business Economics*, 61(4):1497–1514.
- Apruzzese, M., Bruni, M. E., Musso, S., and Perboli, G. (2023). 5G and Companion Technologies as a Boost in New Business Models for Logistics and Supply Chain. *Sustainability*, 15(15):11846.
- Aryee, R., Kanda, W., Geissdoerfer, M., and Kirchherr, J. (2025). Circular ecosystems: Past, present, and future research directions. *Journal of Industrial Ecology*.
- Assenova, V. A. (2018). Modeling the diffusion of complex innovations as a process of opinion formation through social networks. *PloS one*, 13(5):e0196699.
- Azmat, F., Lim, W. M., Moyeen, A., Voola, R., and Gupta, G. (2023). Convergence of business, innovation, and sustainability at the tipping point of the sustainable development goals. *Journal of Business Research*, 167:114170.
- Babovsky, H. and Illner, R. (1989). A convergence proof for nanbu's simulation method for the full boltzmann equation. *SIAM journal on numerical analysis*, 26(1):45–65.
- Bass, F. M. (1969). A new product growth for model consumer durables. *Management science*, 15(5):215–227.

- Bass, F. M. (2004). The bass model. *Manag. Sci.*, 50(12 Supplement):1833–1840.
- Bass, F. M. and Bultez, A. V. (1982). A note on optimal strategic pricing of technological innovations. *Marketing Science*, 1(4):371–378.
- Bass, F. M., Jain, D., Krishnan, T., et al. (2000). Modeling the marketing-mix influence in new-product diffusion. *New-product diffusion models*, pages 99–122.
- Bass, F. M., Krishnan, T. V., and Jain, D. C. (1994). Why the bass model fits without decision variables. *Marketing science*, 13(3):203–223.
- Bastani, M. S., Asadi, S., and Anumba, C. J. (2016). Application of bass diffusion theory to simulate the impact of feedback and word of mouth on occupants' behavior in commercial buildings: An agent-based approach. *Journal of Architectural Engineering*, 22(4):04016013.
- Bastarrianto, F. F., Hancock, T. O., Choudhury, C. F., and Manley, E. (2023). Agent-based models in urban transportation: review, challenges, and opportunities. *European Transport Research Review*, 15(1):19.
- Bayes, T. (1958). An essay towards solving a problem in the doctrine of chances. *Biometrika*, 45(3-4):296–315.
- Bell, D. R. and Song, S. (2007). Neighborhood effects and trial on the internet: Evidence from online grocery retailing. *Quantitative Marketing and Economics*, 5:361–400.
- Berry, E. M., Burlingame, B., and le Coutre, J. (2024). Partnerships for the sustainable development goals: a call for more science. *Frontiers in Nutrition*, 11:1347593.
- Boesch, P. M., Ciari, F., and Axhausen, K. W. (2016). Autonomous vehicle fleet sizes required to serve different levels of demand. *Transportation Research Record*, 2542(1):111–119.
- Boltzmann, L. (1872). Weitere studien über das wärmeleichgewicht unter gasmolekülen. sitzungberichte der kaiserlichen akademie der wissenschaften. mathematisch-naturwissenschaftliche classe 66, 275–370. *The kinetic theory of gases. History of modern physical sciences*, 1:262–349.
- Brás, G. R. and Robaina, M. (2024). National innovation systems and sustainable environmental performance: A cross country analysis. *Environmental Challenges*, 16:100978.
- Breslin, D., Kask, J., Schlaile, M., and Abatecola, G. (2021). Developing a coevolutionary account of innovation ecosystems. *Industrial Marketing Management*, 98:59–68.
- Bronnenberg, B. J. and Mela, C. F. (2004). Market roll-out and retailer adoption for new brands. *Marketing Science*, 23(4):500–518.

- Brotcorne, L., Perboli, G., Rosano, M., and Wei, Q. (2019). A managerial analysis of urban parcel delivery: A lean business approach. *Sustainability (Switzerland)*, 11(12).
- Campani, M. and Vaglio, R. (2015). A simple interpretation of the growth of scientific/technological research impact leading to hype-type evolution curves. *Scientometrics*, 103(1):75–83.
- Cercignani, C. (1988). The boltzmann equation. In *The Boltzmann equation and its applications*, pages 40–103. Springer, New York.
- Chaab, J., Salhab, R., and Zaccour, G. (2022). Dynamic pricing and advertising in the presence of strategic consumers and social contagion: A mean-field game approach. *Omega*, 109:102606.
- Chaparro-Banegas, N., Ibañez Escribano, A. M., Mas-Tur, A., and Roig-Tierno, N. (2024). Innovation facilitators and sustainable development: a country comparative approach. *Environment, Development and Sustainability*, 26(4):8467–8495.
- Chenavaz, R., Paraschiv, C., and Turinici, G. (2021). Dynamic pricing of new products in competitive markets: A mean-field game approach. *Dynamic Games and Applications*, 11(3):463–490.
- Collste, D., Pedercini, M., and Cornell, S. E. (2017). Policy coherence to achieve the sdgs: using integrated simulation models to assess effective policies. *Sustainability science*, 12(6):921–931.
- Cosguner, K. and Seetharaman, P. (2022). Dynamic pricing for new products using a utility-based generalization of the bass diffusion model. *Management Science*, 68(3):1904–1922.
- Crevier, D. (1993). *AI: the tumultuous history of the search for artificial intelligence*. Basic Books, Inc., USA.
- de Cuveland, K., Choi, E., and Shin, M. (2025). Social dynamics of hyperbole—social media sentiment analysis for hype detection in emerging technologies: a conceptual approach. *Technology Analysis & Strategic Management*, pages 1–15.
- Dedehayir, O. and Steinert, M. (2016). The hype cycle model: A review and future directions. *Technological Forecasting and Social Change*, 108:28–41.
- Dehghanimadvar, M., Shirmohammadi, R., Ahmadi, F., Aslani, A., and Khalilpour, K. R. (2022). Mapping the development of various solar thermal technologies with hype cycle analysis. *Sustainable Energy Technologies and Assessments*, 53:102615.
- Della Marca, R., Loy, N., and Tosin, A. (2021). An sir-like kinetic model tracking individuals' viral load. *arXiv preprint arXiv:2106.14480*.
- Dockner, E. and Jørgensen, S. (1988). Optimal pricing strategies for new products in dynamic oligopolies. *Marketing Science*, 7(4):315–334.

- Dosi, G. (1982). Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change. *Research policy*, 11(3):147–162.
- Düring, B. and Wright, O. (2022). On a kinetic opinion formation model for pre-election polling. *Philosophical Transactions of the Royal Society A*, 380(2224):20210154.
- Dyer, J., Cannon, P., Farmer, J. D., and Schmon, S. (2022). Black-box bayesian inference for economic agent-based models. *arXiv preprint arXiv:2202.00625*.
- Fagnant, D. J. and Kockelman, K. M. (2014). The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies*, 40:1–13.
- Fan, Z.-P., Che, Y.-J., and Chen, Z.-Y. (2017). Product sales forecasting using online reviews and historical sales data: A method combining the bass model and sentiment analysis. *Journal of business research*, 74:90–100.
- Fenn, J. and Raskino, M. (2008). *Mastering the hype cycle: how to choose the right innovation at the right time*. Harvard Business Press.
- Fenn, J. and Time, M. (2007). Understanding gartner’s hype cycles, 2007. *Gartner ID G*, 144727.
- Ferrari, A. G., Jugend, D., Armellini, F., Barbalho, S. C. M., and de Carvalho, M. M. (2023). Crossing actors’ boundaries towards circular ecosystems in the organic food sector: Facing the challenges in an emerging economy context. *Journal of Cleaner Production*, 407:137093.
- Fibich, G. and Gibori, R. (2010). Aggregate diffusion dynamics in agent-based models with a spatial structure. *Operations Research*, 58(5):1450–1468.
- Foster, R. (1986). *Innovation: The Attacker’s Advantage*. Summit Books, NY.
- Fruchter, G. E. and Van den Bulte, C. (2011). Why the generalized bass model leads to odd optimal advertising policies. *International Journal of Research in Marketing*, 28(3):218–230.
- Fu, B., Zhang, J., Wang, S., and Zhao, W. (2020). Classification–coordination–collaboration: a systems approach for advancing sustainable development goals. *National Science Review*, 7(5):838–840.
- Gifford, E., McKelvey, M., and Saemundsson, R. (2021). The evolution of knowledge-intensive innovation ecosystems: Co-evolving entrepreneurial activity and innovation policy in the west swedish maritime system. *Industry and Innovation*, 28(5):651–676.
- Goldenberg, J., Libai, B., Solomon, S., Jan, N., and Stauffer, D. (2000). Marketing percolation. *Physica A: statistical mechanics and its applications*, 284(1-4):335–347.

- Grama-Vigouroux, S., Saidi, S., Uvarova, I., Cirule, I., and Sellami, M. (2023). Drivers and barriers of national innovation ecosystems for implementing sustainable development goals: A latvian case study. *IEEE transactions on engineering management*, 71:4188–4204.
- Grazzini, J., Richiardi, M. G., and Tsionas, M. (2017). Bayesian estimation of agent-based models. *Journal of Economic Dynamics and Control*, 77:26–47.
- Gyimah, S., Owusu-Manu, D.-G., Edwards, D. J., Buerthey, J. I. T., and Danso, A. K. (2025). Exploring the contributions of circular business models towards the transition of green economy in the ghanaian construction industry. *Smart and Sustainable Built Environment*, 14(3):859–880.
- Han, J., Ruan, Y., Wang, Y., and Zhou, H. (2021). Toward a complex adaptive system: The case of the zhongguancun entrepreneurship ecosystem. *Journal of business research*, 128:537–550.
- Han, Z., Tang, Z., and He, B. (2022). Improved bass model for predicting the popularity of product information posted on microblogs. *Technological Forecasting and Social Change*, 176:121458.
- Hashemi, F., Gallay, O., and Hongler, M. (2018). Hype cycle dynamics: Microscopic modeling and detection. *arXiv preprint arXiv:1809.08939*.
- Hassani, S. and Hassani, S. (2009). Dirac delta function. *Mathematical Methods: For Students of Physics and Related Fields*, pages 139–170.
- Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109.
- Hesselink, L. X. and Chappin, E. J. (2019). Adoption of energy efficient technologies by households—barriers, policies and agent-based modelling studies. *Renewable and Sustainable Energy Reviews*, 99:29–41.
- Ho, T.-H., Savin, S., and Terwiesch, C. (2002). Managing demand and sales dynamics in new product diffusion under supply constraint. *Management science*, 48(2):187–206.
- Huang, Y., Kang, A., Jiang, Z., Li, J., Ma, Y., Gan, X., and Zhou, B. (2024). An approach based on ecosystem services for assessing progress towards sustainable development goals at both national and provincial levels in china. *Scientific Reports*, 14(1):29250.
- Husainy, A. S., Mangave, S. S., Ingale, A. S., Patil, Y. R., and Koganole, M. (2024). Innovation ecosystems and green building techniques for a sustainable future: Leveraging advanced technologies. *The Asian Review of Civil Engineering*, 13(2):1–10.
- Iizuka, M. and Hane, G. (2021). Towards attaining the sdgs: cases of disruptive and inclusive innovations. *Innovation and Development*, 11(2-3):343–364.

- Islam, H. (2025). Nexus of economic, social, and environmental factors on sustainable development goals: The moderating role of technological advancement and green innovation. *Innovation and Green Development*, 4(1):100183.
- Iyengar, R., Van den Bulte, C., and Valente, T. W. (2011). Opinion leadership and social contagion in new product diffusion. *Marketing science*, 30(2):195–212.
- Jan, S. Q., Junfeng, J., Iqbal, M. B., Raza, A., Naz, M., and Bhatt, T. K. (2025). The impact of entrepreneurial ecosystems and sustainable digital innovation on business performance: a study of gilgit-baltistan pakistan. *Frontiers in Sustainability*, 6:1485680.
- Jarvenpaa, H. and Makinen, S. (2008a). Empirically detecting the hype cycle with the life cycle indicators: An exploratory analysis of three technologies. In *2008 IEEE International Conference on Industrial Engineering and Engineering Management*, pages 12–16. Ieee.
- Jarvenpaa, H. M. and Makinen, S. J. (2008b). An empirical study of the existence of the hype cycle: A case of dvd technology. In *2008 IEEE International Engineering Management Conference*, pages 1–5. IEEE.
- Jiang, Z., Bass, F. M., and Bass, P. I. (2006). Virtual bass model and the left-hand data-truncation bias in diffusion of innovation studies. *International Journal of Research in Marketing*, 23(1):93–106.
- Jiang, Z. and Jain, D. C. (2012). A generalized norton–bass model for multigeneration diffusion. *Management Science*, 58(10):1887–1897.
- Jun, S.-P. (2012a). A comparative study of hype cycles among actors within the socio-technical system: With a focus on the case study of hybrid cars. *Technological Forecasting and Social Change*, 79(8):1413–1430.
- Jun, S.-P. (2012b). An empirical study of users' hype cycle based on search traffic: the case study on hybrid cars. *Scientometrics*, 91(1):81–99.
- Jun, S.-P., Yeom, J., and Son, J.-K. (2014). A study of the method using search traffic to analyze new technology adoption. *Technological Forecasting and Social Change*, 81:82–95.
- Jütting, M. (2022). Crafting mission-oriented innovation ecosystems: Strategic levers for directing collaborative innovation toward the grand challenges. *IEEE Transactions on Engineering Management*, 71:12053–12067.
- Kalish, S. and Lilien, G. L. (1983). Optimal price subsidy policy for accelerating the diffusion of innovation. *Marketing Science*, 2(4):407–420.
- Kangur, A., Jager, W., Verbrugge, R., and Bockarjova, M. (2017). An agent-based model for diffusion of electric vehicles. *Journal of Environmental Psychology*, 52:166–182.

- Keith, D. R., Sterman, J. D., and Struben, J. (2017). Supply constraints and waitlists in new product diffusion. *System Dynamics Review*, 33(3-4):254–279.
- Khodayari, M. and Aslani, A. (2018). Analysis of the energy storage technology using hype cycle approach. *Sustainable Energy Technologies and Assessments*, 25:60–74.
- Konrad, K. (2006). The social dynamics of expectations: The interaction of collective and actor-specific expectations on electronic commerce and interactive television. *Technology analysis & strategic management*, 18(3-4):429–444.
- Konrad, K., Markard, J., Ruef, A., and Truffer, B. (2012). Strategic responses to fuel cell hype and disappointment. *Technological Forecasting and Social Change*, 79(6):1084–1098.
- Kotler, P., Keller, K. L., Ancarani, F., and Costabile, M. (2014). *Marketing management 14/e*. pearson, Harlow, UK.
- Krishnan, T. V., Bass, F. M., and Jain, D. C. (1999). Optimal pricing strategy for new products. *Management science*, 45(12):1650–1663.
- Krishnan, T. V. and Jain, D. C. (2006). Optimal dynamic advertising policy for new products. *Management Science*, 52(12):1957–1969.
- Krook, M. and Wu, T. T. (1977). Exact solutions of the boltzmann equation. *The Physics of Fluids*, 20(10):1589–1595.
- Kumar, S. and Swaminathan, J. M. (2003). Diffusion of innovations under supply constraints. *Operations research*, 51(6):866–879.
- Kupfer, A. (1987). Now, live experts on a floppy disk. *Fortune*, 116(8):69.
- Lafond, F., Bailey, A. G., Bakker, J. D., Rebois, D., Zadourian, R., McSharry, P., and Farmer, J. D. (2018). How well do experience curves predict technological progress? a method for making distributional forecasts. *Technological Forecasting and Social Change*, 128:104–117.
- Lawrence, K. D. and Lawton, W. H. (1981). Applications of diffusion models: some empirical results. *New product forecasting*, pages 529–541.
- Li, S., Jin, Z., et al. (2014). Modeling and analysis of new products diffusion on heterogeneous networks. *Journal of Applied Mathematics*, 2014.
- Li, X., Yuan, J., Shi, Y., Wang, T., Hu, X., Chan, F. T. S., and Ruan, J. (2020). An extended bass model on consumer quantity of b2c commerce platforms. *Electronic Commerce Research*, 20:609–628.
- Liao, H.-T., Pan, C.-L., and Wu, Z. (2024). Digital transformation and innovation and business ecosystems: A bibliometric analysis for conceptual insights and collaborative practices for ecosystem innovation. *International Journal of Innovation Studies*, 8(4):406–431.

- Lighthill, J. (1973). Artificial intelligence: A general survey. In *Artificial Intelligence: A paper symposium*. Science Research Council, UK.
- Linden, A., Fenn, J., et al. (2003). Understanding gartner's hype cycles. *Strategic Analysis Report N° R-20-1971*. Gartner, Inc, 88:1423.
- Liu, H., Li, W., Jiang, J., and Zhang, Q. (2025). Regional innovation ecosystem and local firms' outward foreign direct investment: An empirical study based on symbiosis perspective. *Technology Analysis & Strategic Management*, 37(9):1045–1060.
- Liu, Z. and Stephens, V. (2019). Exploring innovation ecosystem from the perspective of sustainability: Towards a conceptual framework. *Journal of Open Innovation: Technology, Market, and Complexity*, 5(3):48.
- Lokhandwala, M. and Cai, H. (2018). Dynamic ride sharing using traditional taxis and shared autonomous taxis: A case study of nyc. *Transportation Research Part C: Emerging Technologies*, 97:45–60.
- Loy, N. and Tosin, A. (2020). Boltzmann-type equations for multi-agent systems with label switching. *arXiv preprint arXiv:2006.15550*.
- Loy, N. and Tosin, A. (2021). A viral load-based model for epidemic spread on spatial networks. *arXiv preprint arXiv:2104.12107*.
- Marinelli, L., Bartoloni, S., Pascucci, F., Gregori, G. L., and Farina Briamonte, M. (2023). Genesis of an innovation-based entrepreneurial ecosystem: exploring the role of intellectual capital. *Journal of Intellectual Capital*, 24(1):10–34.
- Marzouk, M., Azab, S., Elshaboury, N., Megahed, A., Metawie, M., Hawary, M. E., Ghaith, D., and Bayoumi, A. (2022). Modeling covid-19 effects on sdgs using system dynamics in egypt. *Environmental Science and Pollution Research*, 29(39):59235–59246.
- Masali, G., Bruni, M. E., and Perboli, G. (2025). The statistical mechanics of innovation diffusion. *Annals of Operations Research*, pages 1–22.
- Masali, G., Morea, D., Giglio, C., Iazzolino, G., Perboli, G., and Bruni, M. E. (2026). Beyond linearity: Quantifying sdg saturation in complex national innovation ecosystems. *Journal of Innovation & Knowledge*, 12:100907.
- Masali, G. and Perboli, G. (2024). Diffusion of connected, communicative and automated mobility systems: an agent-based model. In *2024 IEEE Globecom Workshops (GC Wkshps)*, pages 1–6. IEEE.
- Mukherjee, A., Zaveri, P., Patel, R., Shah, M. T., and Munshi, N. S. (2021). Optimization of microbial fuel cell process using a novel consortium for aromatic hydrocarbon bioremediation and bioelectricity generation. *Journal of Environmental Management*, 298:113546.

- Nagy, B., Farmer, J. D., Bui, Q. M., and Trancik, J. E. (2013). Statistical basis for predicting technological progress. *PloS one*, 8(2):e52669.
- Nanbu, K. (1980). Direct simulation scheme derived from the boltzmann equation. i. monocomponent gases. *Journal of the Physical Society of Japan*, 49(5):2042–2049.
- Neath, A. A. and Cavanaugh, J. E. (2012). The bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(2):199–203.
- Neto, J. R., Figueiredo, C., Gabriel, B. C., and Valente, R. (2024). Factors for innovation ecosystem frameworks: Comprehensive organizational aspects for evolution. *Technological Forecasting and Social Change*, 203:123383.
- Norton, J. A. and Bass, F. M. (1987). A diffusion theory model of adoption and substitution for successive generations of high-technology products. *Management science*, 33(9):1069–1086.
- Nylund, P. A., Brem, A., and Agarwal, N. (2021). Innovation ecosystems for meeting sustainable development goals: The evolving roles of multinational enterprises. *Journal of Cleaner Production*, 281:125329.
- Oskam, I., Bossink, B., and de Man, A.-P. (2021). Valuing value in innovation ecosystems: How cross-sector actors overcome tensions in collaborative sustainable business model development. *Business & society*, 60(5):1059–1091.
- Owen, R. and Vedanthachari, L. (2022). Exploring the role of uk government policy in developing the university entrepreneurial finance ecosystem for cleantech. *IEEE Transactions on Engineering Management*, 70(3):1026–1039.
- Palmer, J., Sorda, G., and Madlener, R. (2015). Modeling the diffusion of residential photovoltaic systems in italy: An agent-based simulation. *Technological Forecasting and Social Change*, 99:106–131.
- Pareschi, L. and Toscani, G. (2013). *Interacting multiagent systems: kinetic equations and Monte Carlo methods*. OUP, Oxford.
- Peñarroya-Farell, M., Miralles, F., and Vaziri, M. (2023). Open and sustainable business model innovation: An intention-based perspective from the spanish cultural firms. *Journal of Open Innovation: Technology, Market, and Complexity*, 9(2):100036.
- Perboli, G. and Rosano, M. (2020). A taxonomic analysis of smart city projects in north america and europe. *Sustainability (Switzerland)*, 12(18).
- Peres, R., Muller, E., and Mahajan, V. (2010). Innovation diffusion and new product growth models: A critical review and research directions. *International journal of research in marketing*, 27(2):91–106.

- Pham, Q. H. and Vu, K. P. (2022). Digitalization in small and medium enterprise: a parsimonious model of digitalization of accounting information for sustainable innovation ecosystem value generation. *Asia Pacific Journal of Innovation and Entrepreneurship*, 16(1):2–37.
- Ponsiglione, C., Cannavacciuolo, L., Primario, S., Quinto, I., and Zollo, G. (2021). The ambiguity of natural language as resource for organizational design: A computational analysis. *Journal of Business Research*, 129:654–665.
- Ponta, L., Puliga, G., Lazzarotti, V., Manzini, R., and Cincotti, S. (2023). To copatent or not to copatent: An agent-based model for firms facing this dilemma. *European Journal of Operational Research*, 306(3):1349–1363.
- Prokop, V., Hajek, P., and Stejskal, J. (2021). Configuration paths to efficient national innovation ecosystems. *Technological Forecasting and Social Change*, 168:120787.
- Qiao, Z., Xu, X., Zou, W., and Huang, Y. (2024). Urban sustainable development goals and ecosystem services: Pathways to achieving coordination. *Land Use Policy*, 146:107317.
- Rachul, C. and Zarzeczny, A. (2012). The rise of neuroskepticism. *International Journal of Law and Psychiatry*, 35(2):77–81.
- Rahmandad, H. and Sterman, J. (2008). Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. *Management science*, 54(5):998–1014.
- Reed, J., Svedberg, P., and Nygren, J. (2025). Enhancing the innovation ecosystem: overcoming challenges to introducing information-driven technologies in health care. *Journal of Medical Internet Research*, 27:e56836.
- Risken, H. (1996). *Fokker-planck equation*. Springer, Berlin, Heidelberg.
- Ritala, P. (2024). Grand challenges and platform ecosystems: Scaling solutions for wicked ecological and societal problems. *Journal of Product Innovation Management*, 41(2):168–183.
- Robinson, B. and Lakhani, C. (1975). Dynamic price models for new-product planning. *Management science*, 21(10):1113–1122.
- Rogers, E. M. (2003). *Diffusion of innovations*. Free Press, New York, NY, 5th edition.
- Rotaris, L. and Scorrano, M. (2023). Insights into peer-to-peer carsharing: Modelling and scenario analysis via a bass diffusion agent-based model. *Research in Transportation Economics*, 97:101261.
- Ryan, B. and Gross, N. C. (1943). The diffusion of hybrid seed corn in two iowa communities. *Rural sociology*, 8(1):15.

- Sachs, J., Lafortune, G., Kroll, C., Fuller, G., and Woelm, F. (2022). Sustainable development report 2022—from crisis to sustainable development: the sdgs as roadmap to 2030 and beyond.
- Sahal, D. (1979). A theory of progress functions. *AIIE transactions*, 11(1):23–29.
- Schiera, D. S., Minuto, F. D., Bottaccioli, L., Borchiellini, R., and Lanzini, A. (2019). Analysis of rooftop photovoltaics diffusion in energy community buildings by a novel gis-and agent-based modeling co-simulation platform. *Ieee Access*, 7:93404–93432.
- Schilling, M. A. and Esmundo, M. (2009). Technology s-curves in renewable energy alternatives: Analysis and implications for industry and government. *Energy policy*, 37(5):1767–1781.
- Shen, W., Duenyas, I., and Kapuscinski, R. (2011). New product diffusion decisions under supply constraints. *Management Science*, 57(10):1802–1810.
- Shi, Y., Wei, Z., and Shahbaz, M. (2023). Analyzing the co-evolutionary dynamics of consumers’ attitudes and green energy technologies based on a triple-helix model. *Renewable and Sustainable Energy Reviews*, 171:113009.
- Silvestrini, P., Amato, U., Vettoliere, A., Silvestrini, S., and Ruggiero, B. (2017). Rate equation leading to hype-type evolution curves: A mathematical approach in view of analysing technology development. *Technological Forecasting and Social Change*, 116:1–12.
- Simon, H. and Sebastian, K.-H. (1987). Diffusion and advertising: the german telephone campaign. *Management science*, 33(4):451–466.
- SINFONICA Consortium (2023). SINFONICA web site. <https://sinfonica.eu/>. Last accessed: 2024/08/03.
- Singh, O., Anand, A., Kapur, P., and Aggrawal, D. (2012). Consumer behaviour-based innovation diffusion modelling using stochastic differential equation incorporating change in adoption rate. *International Journal of Technology Marketing*, 7(4):346–360.
- Sinha, R. K. and Chandrashekar, M. (1992). A split hazard model for analyzing the diffusion of innovations. *Journal of Marketing Research*, 29(1):116–127.
- Song, J. (2025). The intersection of climate technology and public diplomacy: Insights into policy and implementation. *Sustainable Futures*, 9:100465.
- Staicu, D. and Pop, O. (2018). Mapping the interactions between the stakeholders of the circular economy ecosystem applied to the textile and apparel sector in romania. *Management & Marketing*, 13(4):1190–1209.
- Steinert, M. and Leifer, L. (2010). Scrutinizing gartner’s hype cycle approach. In *Picmet 2010 technology management for global economic growth*, pages 1–13. IEEE.

- Stummer, C., Kiesling, E., Günther, M., and Vetschera, R. (2015). Innovation diffusion of repeat purchase products in a competitive market: An agent-based simulation approach. *European Journal of Operational Research*, 245(1):157–167.
- Tarde, G. (1903). The laws of imitation, trans. *EC Parsons*. New York: Henry, Holt.
- Thomas, L. D., Ritala, P., Karhu, K., and Heiskala, M. (2025). Vertical and horizontal complementarities in platform ecosystems. *Innovation*, 27(3):369–393.
- Toscani, G. (2006). Kinetic models of opinion formation. *Communications in mathematical sciences*, 4(3):481–496.
- Turner, B. M. and Van Zandt, T. (2012). A tutorial on approximate bayesian computation. *Journal of Mathematical Psychology*, 56(2):69–85.
- United Nations (2015). Transforming our world: The 2030 agenda for sustainable development. *United Nations*.
- Van Lente, H., Spitters, C., and Peine, A. (2013). Comparing technological hype cycles: Towards a theory. *Technological forecasting and social change*, 80(8):1615–1628.
- Von Toussaint, U. (2011). Bayesian inference in physics. *Reviews of Modern Physics*, 83(3):943–999.
- Wei, Y., Zhong, F., Song, X., and Huang, C. (2025). Achieving sustainable development through national innovation: Evidence from 131 countries. *Environmental Impact Assessment Review*, 115:108046.
- Weisstein, E. W. (2002). Heaviside step function. <https://mathworld.wolfram.com/>.
- Wilensky, U. (2015). *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with Netlogo*. The MIT Press.
- Wonglimpiyarat, J. (2025). Achieving the united nations sustainable development goals—innovation diffusion and business model innovations. *foresight*, 27(1):101–119.
- World Economic Forum (2020). Global competitiveness report special edition 2020: How countries are performing on the road to recovery. section 4: Innovation ecosystem. *World Economic Forum*.
- World Economic Forum (2025). From tech-first to climate-first: Why we must redefine how innovation is measured. *World Economic Forum*.
- Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the aeronautical sciences*, 3(4):122–128.
- Xu, R. and Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Transactions on neural networks*, 16(3):645–678.

- Xu, Y. and Li, Y. (2025). A bibliometric review of innovation ecosystem. *Systems Research and Behavioral Science*, 42(4):978–995.
- Yikun, Z., Leong, L. W., Abu-Rumman, A., Shraah, A. A., Hishan, S. S., et al. (2023). Green growth, governance, and green technology innovation. how effective towards sdgs in g7 countries? *Economic research-Ekonomska istraživanja*, 36(2).
- Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American economic review*, 99(5):1899–1924.
- Yuan, H., Wang, X., Gao, L., Wang, T., Liu, B., Fang, D., and Gao, Y. (2023). Progress towards the sustainable development goals has been slowed by indirect effects of the covid-19 pandemic. *Communications Earth & Environment*, 4(1):184.
- Zhang, N., Lu, Y., Chen, J., and Hwang, B.-G. (2022). An agent-based diffusion model for residential photovoltaic deployment in singapore: Perspective of consumers' behaviour. *Journal of Cleaner Production*, 367:132793.
- Zhang, R.-J., Tai, H.-W., Cao, Z.-X., Wei, C.-C., and Cheng, K.-T. (2024). Green innovation ecosystem evolution: Diffusion of positive green innovation game strategies on complex networks. *Journal of Innovation & Knowledge*, 9(3):100500.

Appendix A

SDG Index Peak and Saturation

Good fit
Bad fit, no S-shape or too fluctuating
Under-estimated by Bass curve
Invalid
Smoothed data (moving average)
Smoothed data (exponential)
M forced to 100

Fig. A.1 Legend.

1	Country	p	q	M	Peak	Saturation	RMSE
2	Afghanistan	0.0000	0.6998	50.9085	2015	2033	0.1670
3	Albania	0.0154	0.2580	71.8303	2010	2028	0.0551
4	Algeria	0.0183	0.2187	72.0433	2010	2031	0.0333
5	Angola	0.0014	0.7962	50.8063	2007	2017	0.0757
6	Argentina	0.0182	0.2425	73.1461	2009	2028	0.0453
7	Armenia	0.0205	0.1440	71.9238	2011	2041	0.0408
8	Australia	0.0232	0.1906	75.9763	2009	2031	0.0626
9	Austria	0.0182	0.1992	82.7561	2011	2033	0.0408
10	Azerbaijan	0.0195	-0.0195	97.2088	Invalid	2000	0.0575
11	Bahrain	0.0000	0.7014	64.1737	2014	2038	0.1623
12	Bangladesh	0.0003	0.4158	63.9388	2017	2038	0.1289
13	Barbados	0.0000	3.8992	66.3164	2013	2013	0.5023
14	Belarus	0.0207	0.1086	78.2896	2012	2050	0.0647
15	Belgium	0.0407	0.1344	80.2107	2006	2031	0.0465
16	Belize	0.0196	0.2201	65.8351	2010	2030	0.0517
17	Benin	0.0117	0.2123	52.0348	2012	2036	0.0597
18	Bhutan	0.0160	0.1967	72.0170	2011	2035	0.0290
19	Bolivia	0.0167	0.1712	69.1018	2012	2038	0.0477
20	Bosnia and Herzegovina	0.0061	0.2495	72.6091	2014	2037	0.0523
21	Botswana	0.0001	0.6575	60.9816	2013	2030	0.0425
22	Brazil	0.0315	0.1310	73.5416	2008	2036	0.0510
23	Brunei Darussalam	0.0007	0.5055	66.1557	2013	2029	0.1728
24	Bulgaria	0.0447	0.2959	73.6325	2005	2017	0.0716
25	Burkina Faso	0.0040	0.3137	54.6639	2013	2033	0.0753
26	Burundi	0.0254	0.1912	54.4735	2009	2030	0.0398
27	Cambodia	0.0027	0.3254	63.6005	2014	2035	0.1209
28	Cameroon	0.0066	0.3085	55.4996	2012	2031	0.0666
29	Canada	0.0234	0.1956	77.9108	2009	2031	0.0374
30	Central African Republic	0.0008	0.5053	38.2999	2013	2008	0.7647
31	Chad	0.0008	0.9443	41.0642	2011	2026	0.3510
32	Chile	0.0062	0.2689	78.3844	2013	2035	0.0555
33	China	0.0025	0.3571	72.9148	2013	2033	0.0222
34	Colombia	0.0349	0.1015	71.6421	2007	2040	0.0586
35	Congo, Dem. Rep.	0.0016	0.3789	50.3804	2014	2033	0.0518
36	Congo, Rep.	0.0011	0.5158	51.3965	2011	2026	0.1601
37	Costa Rica	0.0234	0.1432	74.5305	2010	2039	0.0618
38	Cote d'Ivoire	0.0002	0.4939	58.3372	2016	2035	0.0730
39	Croatia	0.0007	0.5288	77.7624	2012	2028	0.0979
40	Cuba	0.0816	-0.0816	76.4703	Invalid	2100+	0.0819
41	Cyprus	0.0277	0.0790	77.2365	2009	2053	0.0611
42	Czech Republic	0.0283	0.1340	81.6887	2009	2037	0.0311
43	Denmark	0.0012	0.4455	85.0730	2013	2030	0.1782
44	Djibouti	0.0000	0.0457	19835.1862	2194	2000	0.1485
45	Dominican Republic	0.0266	0.1332	72.4844	2010	2038	0.0487
46	East and South Asia	0.0020	0.3555	66.3282	2014	2034	0.0432
47	Eastern Europe and Central Asia	0.0028	0.3063	71.5756	2015	2037	0.0667
48	Ecuador	0.0071	0.3351	72.2962	2011	2028	0.0319
49	Egypt, Arab Rep.	0.0054	0.1197	76.5505	2004	2073	0.0657
50	El Salvador	0.0158	0.2773	69.9140	2009	2027	0.0325
51	Estonia	0.0460	0.0792	82.0924	2004	2038	0.0814
52	Eswatini	0.0209	0.2609	54.6184	2008	2026	0.0525
53	Ethiopia	0.0092	0.1288	66.1318	2019	2059	0.0326
54	Fiji	0.0007	0.4079	72.6598	2015	2035	0.1255
55	Finland	0.0251	0.3026	86.3496	2007	2021	0.0567
56	France	0.0271	0.1357	82.4067	2009	2037	0.0328
57	Gabon	0.0206	-0.0206	78.1239	Invalid	2100+	0.0833
58	Gambia, The	0.0017	0.3610	59.5360	2014	2034	0.1519
59	Georgia	0.0071	0.2538	74.0724	2013	2036	0.0347
60	Germany	0.0280	0.1529	83.0503	2009	2034	0.0455

Fig. A.2 The fitting parameters, peak and saturation times and fitting error for the first third of the analyzed entities. The RMSEs in yellow are between 0.15 and 0.2, the ones in red exceed 0.2.

1	Country	p	q	M	Peak	Saturation	RMSE
61	Ghana	0.0090	0.2427	64.0398	2013	2035	0.0505
62	Greece	0.0111	0.0204	100.0000	2019	2087	0.0554
63	Guatemala	0.0069	0.2597	61.2995	2013	2035	0.0940
64	Guinea	0.0193	0.2005	51.7379	2010	2032	0.0521
65	Guyana	0.0048	0.1905	65.6310	2018	2050	0.0857
66	Haiti	0.0004	0.5458	52.9412	2013	2029	0.0946
67	High-income Countries	0.0239	0.1799	77.7987	2009	2032	0.0229
68	Honduras	0.0140	0.2764	63.1187	2010	2028	0.0641
69	Hungary	0.0179	0.0430	85.3209	2014	2097	0.0608
70	Iceland	0.0037	0.3785	78.6239	2012	2029	0.1015
71	India	0.0000	0.5976	58.8186	2015	2034	0.1810
72	Indonesia	0.0025	0.3139	70.1073	2015	2037	0.0362
73	Iran, Islamic Rep.	0.0022	0.3486	68.6304	2014	2034	0.1019
74	Iraq	0.0136	0.1929	62.6021	2012	2017	0.0879
75	Ireland	0.0357	0.2913	80.3227	2006	2019	0.0507
76	Israel	0.0007	0.4269	73.4250	2014	2034	0.0909
77	Italy	0.0310	0.2097	78.4631	2007	2026	0.0452
78	Jamaica	0.0128	0.2465	68.9508	2011	2031	0.0921
79	Japan	0.0067	0.0144	100.0000	2036	2123	0.0493
80	Jordan	0.0003	0.4011	69.5527	2017	2038	0.1866
81	Kazakhstan	0.0068	0.3597	71.3882	2010	2027	0.0561
82	Kenya	0.0030	0.2883	60.8185	2015	2038	0.1831
83	Korea, Rep.	0.0338	0.1455	78.2431	2008	2032	0.0473
84	Kuwait	0.0019	0.3082	63.4129	2016	2038	0.2869
85	Kyrgyz Republic	0.0063	0.2797	74.0953	2013	2034	0.0678
86	Lao PDR	0.0013	0.4090	62.2188	2014	2032	0.1591
87	Latin America and the Caribbean	0.0257	0.1990	69.9185	2009	2029	0.0354
88	Latvia	0.0229	0.1487	81.4161	2010	2038	0.0377
89	Lebanon	0.0209	0.2487	65.2255	2009	2027	0.0971
90	Lesotho	0.0229	0.1485	56.3868	2010	2038	0.0642
91	Liberia	0.0291	0.1531	50.5810	2009	2033	0.0864
92	Lithuania	0.0518	0.1696	75.2229	2005	2023	0.0532
93	Lower-middle-income Countries	0.0012	0.3831	61.2506	2015	2034	0.0953
94	Low-income Countries	0.0192	0.1589	52.0185	2011	2039	0.0246
95	Luxembourg	0.0429	0.0429	81.2562	Invalid	2100	0.0714
96	Madagascar	0.0052	0.0130	100.0000	2050	2198	0.0513
97	Malawi	0.0135	0.2314	53.6629	2011	2032	0.0497
98	Malaysia	0.0032	0.3331	70.6374	2013	2033	0.0471
99	Maldives	0.0103	0.2311	72.0180	2012	2035	0.0462
100	Mali	0.0135	0.2225	54.8403	2011	2033	0.0446
101	Malta	0.0101	0.1858	78.9169	2014	2042	0.0452
102	Mauritania	0.0173	0.2131	56.4918	2010	2032	0.0478
103	Mauritius	0.0028	0.2145	70.1122	2012	2031	0.1839
104	Mexico	0.0203	0.1283	72.5363	2012	2044	0.0480
105	Middle East and North Africa	0.0125	0.1879	67.8083	2013	2039	0.0204
106	Moldova	0.0315	0.1160	74.9729	2008	2039	0.0380
107	Mongolia	0.0000	0.9284	63.4178	2014	2029	0.2668
108	Montenegro	0.0000	2.6663	66.8666	2008	2000	0.2419
109	Morocco	0.0249	0.1822	69.7182	2009	2031	0.0536
110	Mozambique	0.0246	0.1084	56.2393	2011	2046	0.0407
111	Myanmar	0.0002	0.5337	64.7747	2014	2033	0.0763
112	Namibia	0.0007	0.5371	62.7019	2012	2027	0.0610
113	Nepal	0.0098	0.2609	67.0057	2012	2032	0.0259
114	Netherlands	0.0125	0.2638	80.0322	2011	2030	0.0399
115	New Zealand	0.0000	0.0962	19458.5562	2118	2000	0.2045
116	Nicaragua	0.0253	0.2010	67.7272	2009	2029	0.0411
117	Niger	0.0130	0.2123	53.0794	2012	2035	0.0587
118	Nigeria	0.0141	0.0929	58.6410	2017	2065	0.0673
119	North Macedonia	0.0009	0.4802	71.4693	2013	2029	0.0952
120	Norway	0.0343	0.1172	82.8494	2008	2037	0.0542

Fig. A.3 The fitting parameters, peak and saturation times and fitting error for the second third of the analyzed entities. The RMSEs in yellow are between 0.15 and 0.2, the ones in red exceed 0.2.

1	Country	p	q	M	Peak	Saturation	RMSE
121	Oceania	0.0081	0.3099	54.2227	2009	2039	0.3814
122	OECD members	0.0266	0.1693	77.7033	2009	2032	0.0215
123	Oman	0.0008	0.4614	69.1077	2013	2031	0.0628
124	Pakistan	0.0010	0.4097	59.4725	2014	2033	0.0528
125	Panama	0.0003	0.5011	63.5491	2014	2032	0.1932
126	Papua New Guinea	0.0016	0.2985	53.1516	2017	2034	0.3549
127	Paraguay	0.0150	0.2169	67.7118	2011	2033	0.0608
128	Peru	0.0205	0.1744	73.1948	2010	2035	0.0376
129	Philippines	0.0025	0.2629	67.5621	2017	2043	0.1563
130	Poland	0.0164	0.1963	81.5929	2011	2035	0.0525
131	Portugal	0.0270	0.1354	80.6985	2009	2038	0.0507
132	Qatar	0.0218	0.2015	67.3580	2009	2031	0.0498
133	Rest of World (non-OECD)	0.0025	0.3386	63.9853	2014	2034	0.0513
134	Romania	0.0173	0.2989	77.4543	2009	2024	0.0630
135	Russian Federation	0.0010	0.3833	73.9810	2015	2036	0.1161
136	Rwanda	0.0149	0.2112	60.4870	2011	2034	0.0433
137	Sao Tome and Principe	0.0268	0.3445	59.7679	2006	2019	0.0752
138	Saudi Arabia	0.0002	0.4373	66.8010	2017	2038	0.0782
139	Senegal	0.0187	0.1566	60.5501	2012	2040	0.0346
140	Serbia	0.0012	0.3607	75.9732	2015	2036	0.0817
141	Sierra Leone	0.0016	0.3919	53.3651	2013	2032	0.0579
142	Singapore	0.0493	0.1440	71.8849	2005	2026	0.1018
143	Slovak Republic	0.0267	0.1432	79.5257	2009	2036	0.0265
144	Slovenia	0.0309	0.1806	80.3965	2008	2029	0.0481
145	Small Island Developing States	0.0023	0.3652	65.0884	2013	2032	0.1167
146	Somalia	0.0001	0.5395	45.0035	2015	2034	0.1732
147	South Africa	0.0047	0.2697	64.6538	2014	2037	0.0380
148	South Sudan	-0.0050	-1.0879	38.9621	1995	2000	0.5034
149	Spain	0.0216	0.0455	86.0237	2011	2084	0.0439
150	Sri Lanka	0.0008	0.4225	69.5706	2014	2033	0.1119
151	Sub-Saharan Africa	0.0134	0.1929	54.8214	2012	2038	0.0311
152	Sudan	0.0190	0.1652	50.5136	2011	2038	0.0805
153	Suriname	0.0043	0.2439	72.9186	2006	2044	0.1362
154	Sweden	0.0084	0.6389	85.1677	2006	2015	0.0815
155	Switzerland	0.0094	0.3084	80.5728	2010	2028	0.0831
156	Syrian Arab Republic	0.0003	0.3946	54.3356	2012	2000	0.8235
157	Tajikistan	0.0189	0.1397	71.0639	2012	2043	0.0800
158	Tanzania	0.0030	0.3404	56.8759	2013	2033	0.1511
159	Thailand	0.0107	0.2137	74.9094	2013	2037	0.0244
160	Togo	0.0001	0.5485	54.3981	2016	2035	0.2039
161	Trinidad and Tobago	0.0227	0.2400	65.5566	2008	2026	0.0519
162	Tunisia	0.0269	0.1268	71.9330	2010	2039	0.0365
163	Turkey	0.0080	0.4396	69.5037	2008	2021	0.0840
164	Turkmenistan	0.0003	0.5003	64.3003	2014	2033	0.2288
165	Uganda	0.0096	0.1915	56.7479	2014	2042	0.0516
166	Ukraine	0.0030	0.1017	100.0000	2033	2049	0.0604
167	United Arab Emirates	0.0019	0.3523	68.8291	2014	2034	0.0987
168	United Kingdom	0.0032	0.3662	80.6527	2012	2031	0.1659
169	United States	0.0056	0.4218	74.2381	2010	2024	0.0550
170	Upper-middle-income Countries	0.0053	0.2928	71.9489	2013	2033	0.0317
171	Uruguay	0.0066	0.2684	77.8602	2013	2034	0.0376
172	Uzbekistan	0.0072	0.1173	80.4991	2022	2068	0.0243
173	Venezuela, RB	0.0001	0.5000	58.2483	2018	2000	0.7109
174	Vietnam	0.0109	0.2248	73.4607	2012	2035	0.0356
175	World	0.0034	0.3137	66.4799	2014	2034	0.0330
176	Yemen, Rep.	0.0225	0.6710	52.1119	2004	2012	0.1024
177	Zambia	0.0051	0.3584	54.5576	2011	2028	0.0510
178	Zimbabwe	0.0001	0.8142	57.2453	2011	2024	0.1119
179							
180							

Fig. A.4 The fitting parameters, peak and saturation times and fitting error for the last third of the analyzed entities. The RMSEs in yellow are between 0.15 and 0.2, the ones in red exceed 0.2.

Appendix B

Posterior Distributions

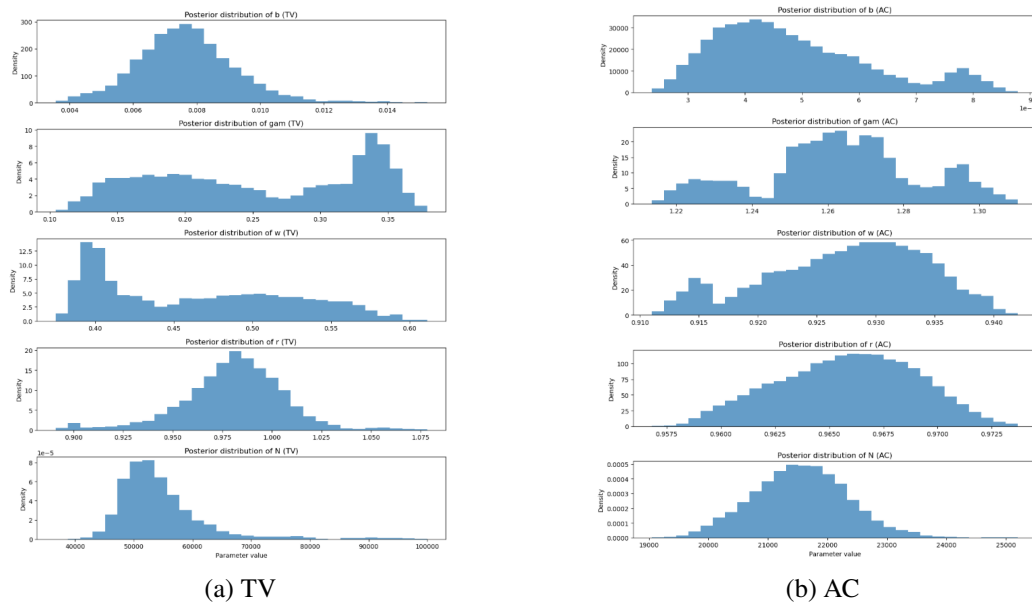


Fig. B.1 The posterior distributions of the parameters obtained through the Metropolis-Hastings algorithm, for the products color TV and air conditioner.

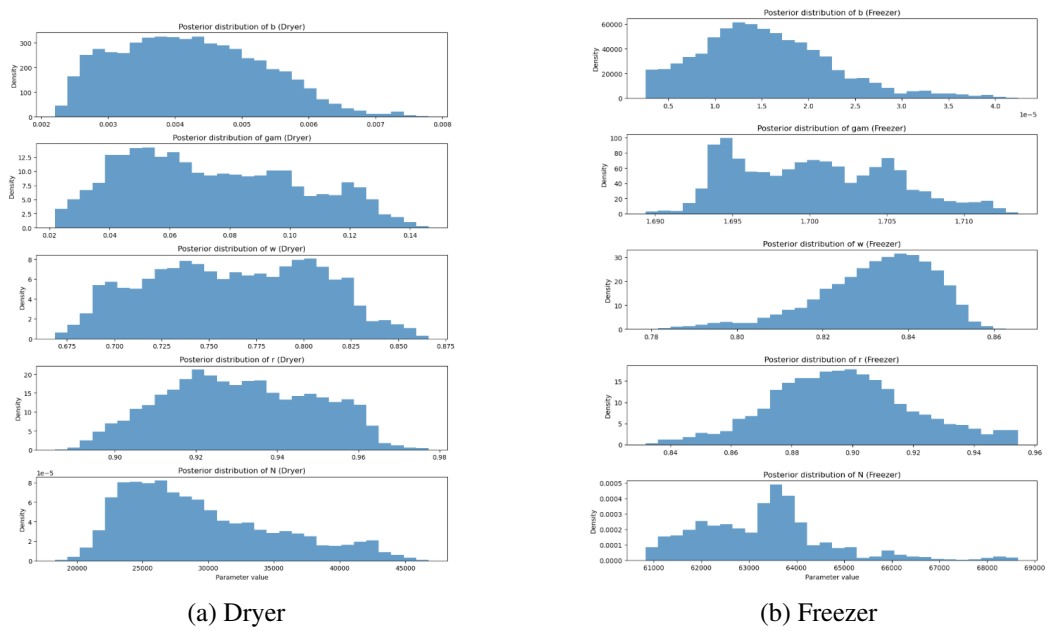


Fig. B.2 The posterior distributions of the parameters obtained through the Metropolis-Hastings algorithm, for the products dryer and freezer.