

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Understanding, measuring, and projecting
technology growth in energy transitions

AVI JAKHMOLA

Department of Environmental and Energy Sciences
Division of Physical Resource Theory
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2026

Understanding, measuring, and projecting technology growth in energy transitions

AVI JAKHMOLA

ISBN 978-91-8103-486-8

DOI: 10.63959/chalmers.dt/5943

ISSN 0346-718X

Doktorsavhandlingar vid Chalmers tekniska högskola, Ny serie 5943

Department of Environmental and Energy Sciences

Division of Physical Resource Theory

Chalmers University of Technology

SE-412 96 Göteborg

Sweden

Telephone: +46 (0)31-772 1000



© 2026 Avi Jakhmola

Pages i–x and 1–90 are licensed under the
Creative Commons Attribution-ShareAlike 4.0 International License.

To view a copy of this license, visit

<https://creativecommons.org/licenses/by-sa/4.0/>.

Paper I is © 2026, Springer Nature Limited.

Papers II and III are published open access articles distributed under Creative Commons licenses.

Papers IV and V are working papers. © 2026, the authors.

Cover: Simulations produced in R by the author.

Chalmers Reproservice

Göteborg, Sweden 2026

Understanding, measuring, and projecting technology growth in energy transitions

AVI JAKHMOLA

Department of Environmental and Energy Sciences
Chalmers University of Technology

ABSTRACT

Meeting climate targets depends on how fast low-carbon technologies replace fossil fuels. Wind and solar power are now among the cheapest sources of new electricity, which many take to mean their growth will be rapid, seamless and uniform. Techno-economic models, in which deployment is largely cost-driven, project much the same. But how does growth actually unfold, and how far ahead can it be anticipated?

This thesis shows that growth instead proceeds in four phases – formative, accelerating, steady, and slowdown – each shaped by drivers and barriers that shift as technology and policy co-evolve. It develops tools that analyse growth as it unfolds, without imposing a fixed shape or destination, and locate shifts between phases directly in deployment data. The findings contradict three common assumptions. Growth does not follow a simple S-curve: for wind and solar, initial acceleration typically ends at around 3% of a country's electricity generation, far earlier than an S-curve's midpoint implies, giving way to prolonged steady growth punctuated by stalls and re-accelerations. Learning rates are not constant: across 145 cost series, the cost decline per doubling of deployment steepens as technologies mature, yet annual cost reductions slow, because each doubling takes longer. And technologies do not spread evenly: across more than 90 technologies in over 200 countries, later adopters pass through the phases in less time but deploy more slowly and saturate lower, consistent with a weaker capacity to absorb new technologies.

A hindcast-validated probabilistic model, PROLONG, projects a baseline for what is likely given the growth observed across countries, rather than what is cost-optimal or required by a target. Its central projections for wind and solar grow much like the IPCC's 2 °C-compatible pathways, and faster than current policy scenarios. The COP28 pledge to triple renewables by 2030 sits near its 95th percentile, and would require major economies to accelerate wind and solar growth by 1.4 to 5 times. Together these findings replace a smooth, cost-driven picture with a phase-based one, in which barriers change from phase to phase rather than fading as technologies become cheap, and the support that sustains growth evolves with them.

Keywords: energy transitions, technology growth, technology diffusion, probabilistic projections, climate change mitigation

APPENDED PUBLICATIONS

This thesis is based on the work contained in the following papers:

- Paper I** J. Jewell, A. Cherp, F. W. Geels, M. Suzuki, L. Nacke, J. Tosun, S. Bhowmik, T. Kazlou, A. Jakhmola and V. Vinichenko (2026). Policy-driven growth of technologies to accelerate climate action. *Nature Reviews Earth & Environment* **7** (4), pp. 235–252.
- Paper II** A. Jakhmola, J. Jewell, V. Vinichenko and A. Cherp (2026a). Rethinking S-curves for policy-driven energy technologies. *Joule* **10**, p. 102526.
- Paper III** A. Jakhmola, J. Jewell, V. Vinichenko and A. Cherp (2026b). Probabilistic projections of global wind and solar power growth based on historical national experience. *Nature Energy* **11** (5), pp. 743–755.
- Paper IV** H. Wang, A. Jakhmola, D. Tong, Q. Zhang, K. He and J. Jewell (2026). The paradox of faster learning and slower cost reduction as technologies mature. *Under review*.
- Paper V** A. Jakhmola and J. Jewell (2026). Declining international support jeopardises deep decarbonisation. *To be submitted*.

Author contributions

Paper I: Conceptualization: J.J. and A.C. Analysis and visualization: M.S., V.V., L.N., S.B., T.K. and A.J. Writing: J.J., A.C., F.W.G., M.S., L.N., T.K., J.T. and A.J. Funding acquisition: J.J., A.C. and J.T.

Paper II: Conceptualization: A.C., A.J., J.J., and V.V. Analysis and visualization: A.J. and V.V. Writing: A.C., A.J., and J.J. Supervision and funding acquisition: J.J. and A.C.

Paper III: Conceptualization: A.C., A.J., J.J. and V.V. Methodology: A.J., A.C., J.J. and V.V. Formal analysis and visualization: A.J. and V.V. Data curation: A.J. and V.V. Writing: A.C., A.J., J.J. and V.V. Supervision and funding acquisition: J.J. and A.C.

Paper IV: Conceptualization: J.J. Methodology: A.J., J.J., H.W. Formal analysis: H.W., A.J. Visualization: H.W., A.J., J.J. Software: H.W., A.J. Data curation: H.W., A.J. Writing – original draft: J.J., A.J., H.W. Writing – review & editing: D.T. Supervision and funding acquisition: J.J., D.T., Q.Z., K.H. (H.W. & A.J. joint first authors).

Paper V: Conceptualization: A.J. and J.J. Formal analysis and data curation: A.J. Visualization: A.J. and J.J. Writing: A.J. and J.J. Supervision and funding acquisition: J.J.

OTHER PUBLICATIONS

Other publications by the author not included in the thesis:

- I. L. Nacke, V. Vinichenko, A. Cherp, A. Jakhmola and J. Jewell (2024). Compensating affected parties necessary for rapid coal phase-out but expensive if extended to major emitters. *Nature Communications* **15** (1), p. 3742.

ACKNOWLEDGEMENTS

Like the technologies I have studied, this PhD moved through several phases, and I owe debts of gratitude to the many people who have been part of each of them.

Jessica Jewell, my supervisor, for giving me this opportunity, for the kind, unwavering support and mentorship through the ups and downs, and for teaching me the importance of asking the right questions. Thank you for always inspiring me to raise the bar. I hope we get to do more fun science together.

Vadim Vinichenko, my co-supervisor, for always being a door-knock away whenever I have felt stuck and needed wise counsel, for your infectious enthusiasm, and for the tangents on books, history, and Bowie.

Aleh Cherp, for MESPOM, for piquing my interest in energy transitions, for your keen-eyed feedback, and for shaping the way I think.

Sonia Yeh, my examiner, and Martin Persson, our director of studies, for your encouragement and for making sure things stayed on course.

Lola and Masa, for being the elder siblings of this PhD (we will finish the full genealogy some day), for never being too busy for a chat, and for being such brilliant collaborators.

ENTAP and the erstwhile POLET group, for the seminars, the research updates, and a camaraderie that cut across Zoom screens; and my co-authors, for making the papers better than I could have alone.

FRT, for being such a vibrant workplace, for the friendly faces in the corridors, and for the lunches, fikas, board-game nights and kick-offs. A special thanks to Angelica and Karin for having the answers to everything, and to everyone who keeps the place running.

Hanna, Andreas, Caroline, and Yodefia, who have been here from the start, and those who joined along the way, for being such fun company and for all the karaokes and after-works.

Ankit, Çağlar, and Ahmet, for making our office a cocoon of warmth and laughter, and for the wide-ranging chats, the meals and the ping-pong.

I must also thank friends and family back home in India and the rest of the world for all their love, and Sofia for drawing me out of my rabbit holes and for the joy you bring. My parents, for supporting me unconditionally every step of the way, and for never discouraging me from asking questions.

/Avi

CONTENTS

Abstract	i
Appended publications	iii
Other publications	iv
Acknowledgements	vii
1 Introduction	1
1.1 Scientific backdrop for anticipating technology growth . . .	3
1.2 Building an empirical basis for analysing technology growth	5
1.3 Aim and research questions	6
1.4 Structure of this thesis	7
2 Background	9
2.1 Energy transitions and climate change mitigation	9
2.2 Technology diffusion, learning, and the S-curve	10
2.3 Socio-technical and political perspectives on technology growth	15
2.4 Projecting technology growth	18
2.5 Research gaps	23
3 Approach and methods	25
3.1 Analytical framework: the four phases of technology growth	25
3.2 Technology and country sample	27
3.3 Model-agnostic tools for measuring technology growth . .	29
3.4 Quantifying cost and learning dynamics	30
3.5 Analysing diffusion between and within countries	32
3.6 Projecting technology growth with PROLONG	33
4 Results	37
4.1 Understanding how technology growth unfolds	37
4.2 Measuring the dynamics of technology growth	46
4.3 Projecting data-driven trajectories for technology growth .	49
5 Discussion	53
5.1 Main contributions	53
5.2 Implications for modelling, policy, and development	54

5.3	Epistemological reflections	57
5.4	Limitations	61
5.5	Future research	62
6	Conclusion	67
	References	71
	Paper I	91
	Paper II	111
	Paper III	129
	Paper IV	155
	Paper V	177

CHAPTER 1

Introduction


“The reader must not suppose that electric lighting by means of [wind] power ... is cheap because the wind costs nothing. On the contrary, the cost of the plant is so great as to more than offset the cheapness of the motive power.”

— Scientific American, on Charles Brush’s wind dynamo, 1890.

When reporting on one of the first attempts to generate electricity using wind power, Scientific American cautioned readers against the seductive appeal of a power source that appeared to cost nothing: while the wind itself was ‘free’, the machine cost a lot of money and was too unreliable to compete with cheap coal power [1]. This view persisted through most of the twentieth century, when the case against wind and solar power rested on their inability to compete not only with fossil energy but also ‘cleaner’ alternatives like hydroelectricity and nuclear power. For more than a hundred years, these relatively expensive technologies remained confined to the fringes of energy systems, even as the world awoke to the perils of climate change and recognised the need to cut greenhouse gas emissions.

That age is firmly behind us. Following rapid technological improvement and burgeoning demand in recent decades, wind and especially solar PV are among the cheapest sources of new electricity in many parts of the world [2, 3]. The very technology costs that prevented them from breaking out of small niches are now used to argue in favour of an unstoppable green energy transition [2–6].

Yet deployment has not been as effortless as that story implies. Despite breaking new records every year, the world is not on course to triple renewables capacity by 2030, nor to reach net-zero emissions by 2050 [3, 7], with countries diverging sharply in how significant a



1. Windmill in the park. 2. Vertical section of the tower. 3. Dynamo. 4. Storage batteries. 5. Regulating apparatus.

Figure 1.1: Charles F. Brush's wind dynamo, Cleveland, 1888. The 17-metre wheel drove a 12 kW dynamo which charged 408 battery cells in the basement of Brush's house and supplied 350 incandescent lamps. Engraving from Scientific American [1]

role new renewables (solar and wind) play in their respective energy systems. For instance, in 2025, solar supplied close to a fifth of all electricity in both Germany and the Netherlands – not the sunniest places in the world – while it barely made a dent in Egypt and Algeria which enjoy some of the best solar resources on the planet [2]. The panels themselves are a globally traded commodity that cost roughly the same in Cairo and Cologne, but there are other economic and socio-political forces that have shaped the expansion of solar power across countries. Thus, the decline in the cost of solar panels or wind turbines, as remarkable a story as it has been, is only one chapter in the energy story of our times.

It is what comes next that motivates this thesis. The bottleneck has shifted from technology invention to expansion: the technologies exist and are cheap, and the central question is no longer whether we have the technologies to decarbonise energy, but how fast, and how widely, they can be deployed. Especially so, given the threat posed by climate change. Mitigating it requires a sweeping technological transformation: the replacement of fossil energy with low-carbon technologies across power, transport, and industry [8]. Meeting the climate targets therefore depends on the *rate* of deployment, and that rate is vexingly hard to anticipate. It is shaped not only by market forces but by policy, and by financial, socio-political, and institutional support. Understanding the patterns of technology growth, and the degree to which its speed can be foreseen, is thus critical to climate mitigation: it conditions the design of policy, the feasibility of targets, and the credibility of different visions for the future.

1.1 Scientific backdrop for anticipating technology growth

If the speed of deployment is not a simple function of technology cost, then there is a significant blind spot in integrated assessment and energy system optimisation models (IAMs and ESOMs), which are the dominant tools we use for anticipating technology futures and exploring how the energy transition might evolve. IAMs and ESOMs represent the growth of technologies largely through cost-optimal dynamics: deployment expands where and when it minimises the cost of the system at large, subject to constraints that are mainly technical or economic in nature [9]. Policy enters this picture only obliquely, most

often as an exogenous input – a carbon price, technology mandate, or emission constraint – rather than as a force that co-evolves with the technology itself. The socio-political drivers that shape real-world deployment, and the way those mechanisms shift as a technology matures are represented scantily if at all [10–12].

The case of solar power demonstrates these limitations plainly. For years, the models in this paradigm were too pessimistic about the prospects of solar power, projecting slow uptake because they assumed solar PV would remain too expensive to be widely adopted. The price of modules collapsed and deployment broke successive records, leaving those projections behind and attracting sustained criticism [6, 13, 14]. The risk now runs in the opposite direction. A technology that has become cheap enough to outcompete all other sources can, in models where cost drives deployment, be projected to expand with insufficient regard for the systemic, financial, and socio-political barriers that govern how fast it can actually be built. The same suite of tools that once projected deployment that was too slow can now project too fast [15, 16], not for any lack of rigour, but because the cost-driven approach is blind to much of what shapes diffusion.

A second family of tools directly extrapolates growth from historical deployment. For decades, researchers have measured the growth of technologies with S-curves, in which deployment accelerates, slows and levels off at a ceiling [17–20]. The S-curve has proven remarkably powerful as a theoretical lens through which to look at the past: one can fit a curve to a technology whose growth is complete and measure how long the transition took, even when the precise functional form remains uncertain [21]. However, it is far less reliable as a guide to the future [22–24], because the ceiling that anchors the curve cannot be known until growth is nearly over [25, 26]. Without the ceiling, no one can say where a still-growing technology is on its S-curve, how long its growth will last, or where it will stop. Nor, while growth is still under way, can anyone be sure that an S-curve describes it at all. Recent work has advanced the frontier by testing which curves best fit the data, and quantifying, through hindcasting, how wrong extrapolations based on such fits have typically been [27, 28]. These methods show how uncertain projections are if growth follows one S-curve or another, but they retain the core assumption that it follows an S-curve.

Both approaches, in different ways, treat growth as a single process. Yet technology growth unfolds within systems shaped simultaneously

by technological, economic, social, and political forces, and the balance between these forces changes over the course of a technology's diffusion [15]. Anticipating how new technologies grow cannot, therefore, rely on a single mechanism operating identically across the whole lifecycle, nor on a single curve.

This thesis advances an understanding of technology growth which bridges techno-economic and socio-political forces without privileging one over the other. It structures growth as a sequence of distinct phases, each governed by a different configuration of drivers and barriers, and each shaped by the co-evolution of technology and policy. Because each phase produces a distinct growth pattern, the transitions between them, such as takeoff and the end of initial acceleration, can be detected in the deployment data itself. That opens a different question. Rather than asking where a technology is on a curve whose shape and ceiling must be assumed, I ask what the data, as observed, tell us about the phase it is in, and develop tools to answer it. Knowing which phase a technology is in matters for two reasons: the policies that can accelerate its growth differ from one phase to the next, and so do the dynamics we can reasonably extrapolate into the future. This phase-based understanding of growth makes it possible to characterise how technologies actually grow and, in turn, measure and project that growth.

1.2 Building an empirical basis for analysing technology growth

Reading phases from the data, rather than from an assumed curve, requires deployment data of a scope that could only recently be assembled. This thesis draws on long-run records of deployment for more than 90 technologies across over 200 countries, from the nineteenth to the twenty-first centuries – not only low-carbon energy technologies such as solar and wind, but a diverse set across sectors and eras, from nuclear power and gas turbines to mobile phones. Many of them have completed their growth, and so offer something the still-unfolding energy transition cannot: the full sequence of phases, from first deployment to saturation. To these I add an expanded dataset linking cost and cumulative deployment for over 60 technologies, which lets me ask whether costs fall at the same pace in every phase of growth.

Breadth of this kind changes what we can ask of the evidence we have. One can fit almost any curve to the history of a single technology retrospectively, but a single case cannot distinguish a recurring pattern from one that happened by chance. Only across many technologies, many countries, and the full course of their growth do robust regularities, or ‘stylised facts’, become visible. Completed cases show what each phase looks like, and the breadth of the sample shows which of those patterns hold across contexts. This is where the aim and questions of this thesis begin.

1.3 Aim and research questions

This thesis aims to develop a better understanding of how technologies grow, to build methods that measure this growth without presupposing its shape, and to use them to project the future deployment of low-carbon energy technologies and benchmark climate-mitigation pathways against probabilistic baselines.

These three ambitions – to understand, to measure, and to project – organise the thesis and its questions. Together they answer a single overarching question:

What characterises the growth of technologies across their lifecycle, and what can an empirically grounded understanding of it offer for measuring and projecting technology growth in energy transitions?

I break this down into three research questions.

RQ1: Understanding growth. What patterns characterise the growth of technologies: (a) how well do the smooth S-curves conventionally used to describe deployment capture how it actually unfolds over the lifecycle; (b) how do costs and learning dynamics evolve over the technology lifecycle; and (c) how does diffusion differ between early and late adopters, both in how quickly a technology reaches new countries and in how it grows within each one?

RQ2: Measuring growth. What tools and metrics can identify phases, turning points, and the speed of diffusion without assuming a functional form in advance?

RQ3: Projecting growth. What methods can probabilistically project the global growth of technologies from past data, and how do the resulting trajectories compare with climate-mitigation pathways?

1.4 Structure of this thesis

The remainder of the thesis proceeds as follows. Chapter 2 situates the work within the study of energy transitions and technology diffusion, drawing on the S-curve, socio-technical, and techno-economic traditions of analysing technology growth to distil the gaps it addresses. Chapter 3 sets out the analytical framework – the four phases of technology growth – together with the data and the model-agnostic tools developed to measure growth, and the probabilistic model developed to project it. Chapter 4 presents the results along the three lines set out above: understanding how growth unfolds, measuring its dynamics, and projecting future trajectories. Chapter 5 draws out the thesis's contributions and their implications for modelling and policy, before turning to its limitations and to directions for future research. Chapter 6 concludes.

CHAPTER 2

Background

2.1 Energy transitions and climate change mitigation

Energy has always been fundamental to human society. The availability of abundant and useful energy underpinned industrialisation, and with it the growth in economic output and living standards that followed [29, 30]. Access to energy remains closely entwined with human development, and its role in modern society is evident from the sheer scale of the systems that supply it.

Those systems have been in flux for much of history. Human societies moved from traditional biomass to coal, and later from coal to oil and gas, and each of these transitions reshaped economies, settlements and ways of life [30–32]. Each also unfolded over decades, rather than years, driven largely by the superior cost, convenience or performance of the incoming energy source [30, 33, 34].

A further transition is now under way, from fossil fuels toward low-carbon electricity led by wind and solar power. It differs from its predecessors in two important respects. First, it is being deliberately driven and sustained by public policy rather than emerging from cost and convenience alone. Second, its pace is dictated by an urgent deadline.

That urgency comes from the need to mitigate climate change (Figure 2.1). Fossil energy produces most of the world’s greenhouse gases [8], so no country can cut emissions far without decarbonising how it generates and uses energy. This is reflected in government commitments including the Paris Agreement, a growing set of national net-zero pledges, and the COP28 promise to triple global renewable capacity by 2030 [3, 7, 8, 35].




Figure 2.1: Total global emissions (a) and electricity generation (b) in the International Energy Agency's Net Zero Emissions by 2050 scenario [3].

Designing climate policy, and anticipating what it will achieve, therefore requires an understanding of how technologies grow. In the rest of this chapter I examine three traditions, each of which has taken up one of three questions: what growth looks like, what drives it, and how it can be projected. The first, rooted in diffusion research and economics, measures how technologies spread and how their costs fall (section 2.2). The second, rooted in the socio-technical and political study of transitions, explains what makes that spread possible (section 2.3). The third, techno-economic modelling, draws on elements of both to look forward (section 2.4). The first two study growth as it has happened; the third anticipates it. Section 2.5 sets out the three gaps these traditions leave, in what is known about growth, in how it is measured, and in how it is projected, and those gaps frame the research questions of this thesis.

2.2 Technology diffusion, learning, and the S-curve

Researchers have studied how technologies spread for more than a century, first as a sociological question and later as an economic one. The earliest scholarly study of diffusion is usually credited to Gabriel Tarde, whose *Les lois de l'imitation* (1890) described the spread of social norms, laws, and customs through imitation and identified the S-shaped curve of the adoption of an innovation [36]. One of the

seminal empirical studies, however, is Ryan and Gross' analysis of the uptake of hybrid corn in two Iowa farming communities [37], which advanced a theoretical exploration of diffusion as a social process in addition to formalising the S-curve of adoption over time. Everett Rogers, himself a rural sociologist trained at Iowa State, synthesised and codified this body of work in *Diffusion of Innovations* [38], and it is largely through his formulation that the paradigm entered the wider social sciences.

The economic reading¹ developed in parallel but drew different conclusions from the same evidence. Griliches also examined the hybrid corn case but explained the observed differences in adoption timing by the *profitability* of the innovation to the individual farmer rather than by social structure or communication [17]. That interpretation provoked an intense debate between sociologists and economists which ran for decades [25, 43–45]. Mansfield extended the economic approach from farmers to firms, modelling the rate at which an innovation is imitated across an industry and again finding diffusion speed to be governed largely by the profitability of adoption [46]. The disagreement was never resolved so much as institutionalised. The two interpretations of the same Iowa data, one explaining adoption through social structure and the other through cost, are the seeds of the traditions I describe in sections 2.3 and 2.4.

The economic reading also gave rise to the literature on technological substitution and modelling which most directly informs the study of energy transitions. Fisher and Pry formalised the displacement of one technology by another as a logistic substitution process [47], which Marchetti and Nakicenovic later applied to the evolution of the global primary energy mix [18]. This approach underpins a substantial subsequent literature on the diffusion of innovations and energy technologies emerging from the International Institute for Applied Systems Analysis [19, 20, 31, 33, 34, 48–50]. Studies in this lineage typically represent deployment with an S-shaped curve defined by a

¹This belongs to a much larger economics of technological change, divided broadly between a neoclassical account in which technical change enters growth as an exogenous shift in the production function, later endogenised as learning [39, 40], and an evolutionary tradition built on firm routines, search and selection, and technological trajectories [41, 42]. Its units of analysis are the firm, the industry and the innovation system rather than, as here, the deployment of a specific technology in a specific country. It nonetheless weaves through what follows, most visibly in Arrow's learning by doing and in the evolutionary roots of the socio-technical tradition (section 2.3).




Figure 2.2: The S-curve of technology diffusion. Adapted from Grubler (1998) [33].

ceiling or saturation level², and characterise the speed of technology growth by the time deployment takes to expand from one fraction of the ceiling to another (Figure 2.2).

A transition’s duration is a span of calendar years; it says nothing about how steeply deployment rises, or how high it finally reaches. The distinction is especially important when the same technology is compared across countries, and is at the heart of a longstanding debate on whether latecomers gain from the experience, proven designs, established supply chains and lower costs they inherit [31, 33, 53, 54], or lose from the weaker financial, human and institutional capacity with which they take a technology up [33, 34, 55, 56]. Technologies now reach new countries far sooner than they once did [55–57]. Whether their lifecycle within a country is also shorter once they arrive, and whether deployment there grows as fast and as far, are separate questions, and the standard metrics do not separate the three.

The S-curve also divides growth into stages. At the very basic level,

²The S-curve has a parallel life in marketing, where the Bass model and its descendants forecast new-product sales from an assumed market potential [51, 52]. I do not engage this literature directly, though it contends with many of the same parameter-estimation problems.

its own geometry segments growth into acceleration up to the inflection point and deceleration after it. The diffusion literature additionally delineates a period before technologies begin ascending along the S-curve: a formative phase, lasting from the deployment of the first commercial units to the onset of sustained growth [58]. Measured across a range of energy technologies, that phase typically lasts two to three decades [58, 59]. Transition studies use a similar sequence of predevelopment, takeoff, acceleration and stabilisation [60], and recent work argues that wind and solar have entered the “next phase” in which the challenges of system integration replace those of cost [61].

Two shortcomings of this work matter here. First, the literature has no concept of a middle phase between acceleration and slowdown: once past the inflection point, growth can only decelerate. Second, no formal method locates the shifts between phases. For instance, existing methods to identify the end of the formative phase rely on thresholds sensitive either to assumptions about the final deployment ceiling [58], or the scale of analysis [62], so a still-growing technology cannot be placed reliably in any phase. This is relevant because the mechanisms of later phases cannot be inferred from earlier ones (Paper I). Whether the slower-than-exponential growth of solar PV since the mid-2010s marks a temporary fluctuation or the end of acceleration, for instance, has major implications for the projections of its future. Moreover, explaining a transition, or comparing it across countries, first requires identifying when it begins.

The vocabulary around technology growth and its description as a *diffusion* process is worth pausing on, because it braids together metaphors from three disciplines: physics, biology and anthropology. From physics comes the image of a substance spreading passively down a gradient, driven by nothing but its own properties and those of the medium. The mathematics operationalising this metaphor, though, comes from biology via two different routes. The logistic curve originated in Verhulst’s work on population growth [63], which was rediscovered a century later by Pearl and Reed, who fitted it to human populations and defined it as growth constrained by a carrying capacity [64]. It is in that ecological sense that the curve reaches the substitution literature, where Marchetti would later treat technologies as species filling niches under a fixed overall market [65]. It reaches the economic diffusion literature by a different path, through epidemic models in which technologies spread like contagions, with adopters

‘infecting’ non-adopters [46]. Either way, the iconic S-curve at the heart of the field comes from biology rather than physical diffusion. The third sense comes from anthropology, whose diffusionist schools of the late nineteenth and early twentieth centuries explained cultural traits as emanating outward from centres of origin [66]. Each discipline lends the word one implication. Physical diffusion implies passivity. Biology, by either route, implies inevitability: a population filling its niche and a contagion running its course both proceed to their end once begun. Anthropological diffusionism implies a direction of travel from centre to periphery. I return in section 5.3 to how well these implications hold up to empirical scrutiny.

Whatever the metaphor, this diffusion strand tracks deployment. A second strand, the technological learning literature, tracks costs instead: not how far technologies spread but how their costs fall as they do. Experience or learning curves relate what a technology costs per unit to how much of it has been built, through the power law known as Wright’s law [40, 67]; a parallel formulation makes cost a function of elapsed time rather than of cumulative production [68–70]. Past studies have estimated learning rates across a wide range of energy technologies [71, 72]. Unlike the diffusion metrics, much of this literature explicitly focuses on tracing the mechanisms underlying cost dynamics. It decomposes observed cost declines into distinct components, attributing them to factors such as plant scale, device efficiency, input prices, research and development, and organisational learning [73–75]. It also explains why some technologies exhibit faster learning than others, by their design complexity and unit size [20, 76]. Cost and deployment drive one another in both directions, so the two strands describe a single process from opposite sides, even though they have largely developed apart.

Each strand has its own limitations for analysing energy transitions. The diffusion strand has two. First, it is predominantly retrospective and mechanism-agnostic. It quantifies historical trajectories and compares metrics across technologies with considerable rigour, but generally does not identify and explain the mechanisms that produce the observed outcomes. Put simply, it describes growth more than it explains it. Second, its methods depend on knowing the ceiling. Because the metrics rely on curve-fitting and estimating growth parameters which include the final deployment ceiling, they work best on technologies whose diffusion is already complete [25, 26]. For

a still-growing technology the ceiling is highly uncertain, and so is everything anchored to it: where the technology is on its curve, how long its growth will last, and where it will stop. The data can also be too short to tell which functional form fits, or whether an S-curve fits at all, and standard S-curves cannot reproduce the recurrent pulses that technology deployment often shows (Paper I). Recent work quantifies these uncertainties by comparing candidate curves and hindcasting their extrapolations, a development taken up in section 2.4.

The learning strand faces a different difficulty. Having traced the mechanisms, it has nonetheless treated the empirically observed learning or cost-reduction rate as a stable property of a technology [72, 77], and most energy system models operationalise it as a technology-specific constant [78], even though observed rates vary substantially over time for a single technology [79–81]. Moreover, whether costs fall because experience accumulates or because of other autonomous mechanisms remains hard to separate [70].

Neither strand yields a reliable, forward-looking account of how fast deployment itself proceeds. The learning strand explains what a technology costs, not how quickly a country deploys it, and the diffusion strand has developed largely apart from it. As a result, the relationship between a technology’s position in its lifecycle and the pace at which its costs fall has gone largely unexamined.

2.3 Socio-technical and political perspectives on technology growth

Technological change has also been studied from very different starting points, in a tradition rooted in the sociology and history of technology, evolutionary economics, political science and political economy. Two perspectives within it matter most here, one socio-technical and one political. Where the diffusion strand asks how fast and how far a technology spreads, these perspectives ask what makes its spread possible at all.

The socio-technical perspective treats technological change as a social phenomenon rather than a purely technical one. Frameworks such as the multi-level perspective and the technological innovation systems approach situate a technology within a wider configuration of actors, institutions, infrastructures, user practices and cultural meanings, and explain change through the interaction of niche innovations,

incumbent regimes and broader landscape pressures [82–86]. Applied to energy, this literature has produced empirical findings as well as frameworks: that incumbent regimes resist displacement rather than yielding to superior alternatives [87]; that emerging technologies depend on protected niches in which they can develop before meeting full market selection [82, 88]; and that transitions are contested political processes whose outcomes depend on which actors can mobilise [84, 89]. Seen this way, a technology succeeds not because it is inherently superior to alternatives but because a system reorganises around it.

The political perspective takes the state as its primary unit of analysis and asks how national policy shapes the evolution of energy systems. Drawing on theories of policy learning and policy diffusion, punctuated equilibrium, the advocacy coalition framework and the politics of interest-group mobilisation, it examines how support for a technology is built, sustained, contested and sometimes withdrawn [90–94]. Technologies that depend on policy are exposed to the volatility of politics in a way that technologies driven by cost alone are not [95, 96].

Both perspectives converge on the idea that technologies and the policies that govern them co-evolve [97, 98]. Cherp and colleagues make that convergence the basis of a framework, conceptualising national energy transitions as the co-evolution of three kinds of system: the energy flows and markets studied by techno-economic analysis, the knowledge, practices and networks that constitute energy technologies, and the systems of political action that produce energy policy [90]. Each system has its own literature, and the framework's claim is that none of them can fully explain energy transitions in isolation from the others.

The relationship between technology and policy runs in both directions: policy shapes deployment, and deployment in turn reshapes the politics of support. The feedbacks this creates pull in opposite directions. Some reinforce growth. Deployment lowers costs and improves performance, which strengthens the case for continued support, which in turn enables further deployment [99]. Beyond costs, deployment builds an industry with employment and revenue to protect, creates constituencies with an interest in continuity, and generates the operating experience that makes further support politically defensible [100–102].

Others resist growth. As a technology becomes pervasive it impinges on incumbents, landscapes and existing systems, mobilising opposition among the actors it displaces and sometimes leading to the weakening or withdrawal of the support on which its growth depends [89, 95, 103–105]. For wind and solar, grid integration challenges grow harder as their share of generation rises; land use and visual impact draw local opposition once deployment is dense; and the fiscal cost of support becomes politically contested [105–107]. In politics, these two kinds of feedback map closely onto a distinction historical sociologists draw: reinforcing feedbacks produce *self-reinforcing sequences*, which entrench a trajectory, and resisting feedbacks produce *reactive sequences*, in which each move provokes a countermove that can slow or reverse it [108]. Reactive sequences offer one reason why growth might stall and later restart, rather than rise smoothly towards a ceiling. Which kind of feedback dominates is not fixed but depends on how far the technology has already been deployed [61, 109, 110].

This tradition also thinks in phases. It identifies them by the mechanisms at work [60, 61], which is what makes its accounts explanatory, but cannot measure when one phase ends and another begins. Thus, even though the literature highlights how the challenges a technology faces, and the policies that can help it, change as it moves from one phase to the next, it cannot say when this shift happens for a given case.

For all its explanatory strength, this tradition has a limitation that mirrors the one identified in section 2.2. It relies predominantly on qualitative methods and on detailed case studies of particular technologies in particular places. This yields rich accounts of mechanisms but makes generalisation difficult, and it leaves the insights hard to operationalise in the forward-looking assessments that are required to inform climate policy [10, 111]. Scholars within the tradition disagree over how far this is a flaw. Critics argue that its frameworks organise description rather than supply causal explanation, and that single cases cannot show which mechanisms were necessary to the outcome [112]. Proponents reply that for large, heterogeneous and slowly unfolding change, mid-range frameworks and recurring mechanisms are the appropriate form of generalisation, while conceding that many case studies remain case histories [108]. Either way, where the diffusion strand measures without explaining, this tradition explains without readily generalising.

2.4 Projecting technology growth

Despite all their attention to the past, none of these traditions has ever been merely historical. In studying how technologies grew, each has ultimately sought to say something about the future, and how fast different technologies can grow to meet climate and other sustainability challenges. Anticipating how fast technologies will grow is a practical necessity: for policymakers setting targets and directing resources, investors allocating capital, and bodies like the IPCC that seek to synthesise knowledge and inform debate.

The task of performing assessments of technology futures has been dominated by what Cherp and colleagues call the techno-economic perspective, which treats energy systems as flows and markets and studies them with the tools of energy systems analysis and economics [90]. Its intellectual roots lie in neoclassical economics. Much of that literature treats energy and other scarce resources as flowing toward an equilibrium, an optimal balance of supply and demand. Markets reach that balance through transactions among rational, self-interested agents; models without markets reach it through the decisions of a perfectly efficient and omniscient social planner [9, 113]. Concerns over environmental degradation, resource scarcity and energy security in the 1970s, together with rapid advances in computing, turned this analytical apparatus toward global problems. The *Limits to Growth* study pioneered one of the first global models of the co-evolving future of society and the environment [114], and the oil crisis of 1973–74 lent urgency to the question of how finite fossil resources should be allocated over time. Nordhaus built one of the first energy-economy models in 1973 [115]. Such models spread quickly through initiatives such as the Energy Modeling Forum in the United States and the Energy Project at the International Institute for Applied Systems Analysis, whose work informed the landmark *Energy in a Finite World* [116].

Integrated assessment models (IAMs) grew out of this work, beginning with Nordhaus' coupling of an energy-economy model to a simple emissions model [117, 118]. Over the following two decades the models added more complex climate representations and modules for population and land use, while retaining the neoclassical lens they had inherited [113]. Alongside IAMs, energy system optimisation models (ESOMs) drew on economics and energy engin-

engineering to develop detailed representations of energy flows, and explore optimal configurations to deliver energy services under different resource, technological, and policy constraints [119]. Together IAMs and ESOMs now inform national policy, underpin agreements from Kyoto to Paris, and increasingly shape public and political discourse about the transition [113, 120–122].

Models of this kind distinguish *endogenous* variables, calculated within the model as it runs, from *exogenous* parameters, which are supplied from outside and held fixed [9, 123]. Deployment of a technology is typically endogenous, expanding where and when it minimises system cost subject to constraints that are mainly technical or economic. So, increasingly, is technological learning, though endogenising it adds structural complexity and its effectiveness in reproducing realistic cost dynamics remains contested [6, 70, 78, 124]. Policy, by contrast, generally enters the models exogenously: a carbon price, a technology mandate or an emissions cap is imposed as a condition, and the model computes the cost-optimal response to it. This is an explicit design choice, and it limits the feedback between policy and technology inside the models, even though in reality the two co-evolve in the manner described in section 2.3.

The models' architecture also dictates what can be represented at all. Geophysical processes and techno-economic quantities, such as resource volumes, installed capacities and operating costs, are readily expressed as quantitative parameters and are consequently well represented [9, 10]. Socio-political factors are not: public acceptance, resistance from incumbents and governance capacity have all proven considerably harder to quantify and integrate, though recent work has begun to do so [11, 123, 125–128].

Each such addition requires someone to translate a mechanism into a form that fits the model's architecture and logic. That translation is rarely straightforward. It obliges modellers to assume how the mechanism should be parameterised and how it interacts with everything else inside the model, and so far only a small share of the mechanisms that shape real transitions has been carried across [12, 111, 123, 128]. Scholars disagree over why it is so. One view argues that the barrier is one of principle: theories that explain change through contested meanings, taken-for-granted cultural structures or conflicts of interest rest on assumptions about social reality that the rational-actor foundations of IAMs cannot absorb, so they can be set alongside

models but not built into them [10]. The other holds it to be one of evidence: a social process can be absorbed into models once it has been measured as a generalisable, quantifiable pattern [11]. Three strategies have been tried: *bridging*, in which models and qualitative analysis run in parallel and meet only at defined moments, typically around shared concepts; *iterating*, in which storylines set model inputs and, in some cases, are revised in light of the outputs; and *merging*, in which social mechanisms are written into the model's structure [11, 111]. The three differ in how far they assume that social mechanisms can be modelled at all, and in that sense map onto the disagreement above, though none has been shown to outperform the others [111]. The qualitative storylines of the Special Report on Emissions Scenarios [129], and later the Shared Socioeconomic Pathways [130, 131], are the most widely used form of iterating, and in practice its unidirectional version: they cover only a handful of socio-economic variables, and because they are supplied from outside they interact only weakly with the variables working inside the model [11, 113]. A further consequence follows from the spatial scale: typically calibrated on long-run global or regional trends, these models obscure the spatial unevenness of transitions and the divergent capacities and characteristics of individual countries [15, 16, 132].

Two consequences follow for this thesis. First, these tools tell us what a target *requires*, or what would cost least, but not what is *likely* to happen [133]. That confines them to what-if questions, and their outcomes are sensitive to what one assumes about which technologies will be available, how cheap they become, and how tightly emissions are capped. Those assumptions have driven successive debates over the prospects of solar photovoltaics [6, 13, 134, 135], negative-emissions technologies [136, 137] and, more recently, electrolytic hydrogen [138, 139]. Second, and more consequentially: a model cannot contain every mechanism that shapes real-world change, so showing that a pathway works inside a model does not mean that it will work outside one. Models can rule a pathway out under a given parameterisation, but among pathways in the “solution space” [140, 141] they cannot separate those the world could actually realise from those it could not [120, 142, 143]. The most developed qualitative response to this challenge comes from bridging. Socio-technical scenarios compare model pathways against analyses of recent niche and regime dynamics, identify transition bottlenecks where the two are

in conflict, and write storylines of the social and political change that would allow the modelled pathway to be realised [144, 145]. But the storylines are written to make a modelled pathway plausible rather than to test it, their authors report having to stretch them to fit some model outputs, and they assign no likelihood to one pathway over another. Judging the real-world feasibility of the scenarios these tools generate therefore remains an open methodological problem.

Kahneman and Lovallo distinguish an *inside view*, which reasons forward from the particulars of a case, from an *outside view*, which reasons by analogy from a distribution of outcomes across comparable cases [146]. Jewell and Cherp bring that distinction to the assessment of feasibility in the climate literature [12]. Making models more realistic by importing evidence from the social sciences takes an inside view, and such additions do not necessarily improve realism. Building more complex and more detailed models, they argue, “may paradoxically increase over-confidence in the likelihood of a given storyline and is more likely to solidify the perception of scenarios as realistic, rather than improve their actual realism” [12]. An outside view proceeds differently. Rather than specifying the mechanisms at work in a given case, it asks what has actually happened across many similar cases and treats that distribution as evidence about what is plausible.

This distinction helps position two lines of work, each responding to a different weakness in the modelling landscape. The first reforms the models from within, and two reforms stand out. *Simulation models* target the behavioural assumptions: drawing on the mechanism-rich socio-technical view, complex-systems and agent-based models simulate heterogeneous actors, learning feedbacks and path dependence rather than solving for a least-cost configuration [147, 148]. Some import measured social factors such as public opinion or the cost of capital directly into a model and test how well the model reproduces history [126, 149]. *Near-optimal methods* instead move away from deriving a single, cost-optimal configuration: modelling to generate alternatives and related methods map the space of solutions lying within a small cost margin of the optimum, and expose the trade-offs among them [150–152].

Both reforms work from inside the modelling paradigm, and each has a limit. Simulation models improve the representation of mechanisms and allow growth to be simulated from the bottom up, but their realism is constrained by what modellers can specify in advance:

which actors are relevant, how they are connected, what rules they follow, and which parameters stay fixed while the model runs [153–155]. Testing models augmented with these additions against history illustrates both that they can pay off and what it takes. Across 31 European countries, adding social factors improved how well a power-sector model reproduced 1990–2019 in all but two, but some factors made the fit worse, gains stopped after about four factors, and the combination that best reproduced installed capacity differed from country to country and in eleven countries worsened the fit for emissions or renewable share [126]. Social evidence improves realism when it is chosen separately for each country, and even then, as the authors note, a structure that fits the past offers no guarantee once the structure of the system itself changes. Near-optimal methods explore the solution space far more richly, but how far they explore is itself fixed in advance, by the cost margin the modeller allows. Recent work uses hindcasting to show how consequential that choice is: real transitions in 31 European countries departed from the cost-optimal path by a median of 44 per cent over thirty years, and only one country fell within the 10 per cent margin commonly used in the literature [156]. Ensuring projected futures encapsulate real-world outcomes may thus require setting wider margins, which would make the range of admissible futures very large. Either way the modelled outcomes define an envelope, and do not assign likelihood to different pathways.

The second line of work takes something closer to the outside view, and targets the gap between the required and the likely. Building on the diffusion literature, data-driven approaches extrapolate from observed trajectories, capturing the aggregate outcome of many mechanisms without needing to specify each one [15, 16, 49, 137, 157, 158]. This has the advantage of grounding expectations in what has happened. But because the past is evidence about what is plausible only while the conditions that produced it persist [10], how far an observed pattern can be carried forward is itself an empirical question. Such approaches have also generally inherited the assumption of smooth trajectories from the S-curve tradition, and have typically produced single, deterministic projections. That has begun to change. A recent body of work pushes data-driven projection toward probabilistic, likelihood-based assessment [6, 27, 28, 77, 138, 139, 159], and it is this frontier I set out to advance.

2.5 Research gaps

The study of how technologies spread has come a long way from two Iowa farming communities and a single crop. Along the way the field divided, and its traditions have since developed largely in parallel. Diffusion research measures growth and finds broad patterns, but does not explain what produces them (section 2.2). The socio-technical and political tradition explains the mechanisms, but struggles to generalise from one case to the next (section 2.3). Techno-economic models answer a different question again: what would need to happen to meet a target, or what would cost least, rather than what is likely to happen (section 2.4). Between them, the three traditions leave three gaps, in what is known about growth, in how it is measured, and in how it is projected, and those gaps set the aim of this thesis (section 1.3).

The first gap, in understanding growth, has three parts. The middle of the lifecycle falls between the two traditions that study growth as it happens. The S-curve assumes acceleration is immediately followed by slowdown, and while the socio-technical literature does identify phases it reads them from the mechanisms at work in particular cases, not from the deployment data itself. Neither says anything about the long stretch after growth stops accelerating but before it saturates (sections 2.2, 2.3).

The other two parts come from the field being split. Because the diffusion and learning strands developed separately, no one has examined how the pace of cost decline depends on where a technology sits in its lifecycle (section 2.2). And because the standard metrics mix together three dimensions of diffusion (how quickly a technology reaches new countries, how long it takes to pass through the lifecycle within each one, and how fast and how far each one deploys it), the evidence on whether latecomers gain or lose remains inconclusive (section 2.2). Together these define the first research question: what patterns characterise how technologies grow? Specifically, how well do the smooth S-curves conventionally used to describe deployment capture how it actually unfolds (RQ1a), how do costs and learning dynamics evolve over the lifecycle (RQ1b), and how does diffusion differ between early and late adopters, both in how quickly a technology reaches new countries and in how it grows within each one (RQ1c)?

The second gap, in measurement, follows from the first. Curve-fitting methods work best when the ceiling is already known, leading

to low reliability when used to study technologies which are still growing. Common threshold metrics need that same ceiling, and thus no existing method can reliably locate the boundaries between phases for technologies that are not approaching saturation (section 2.2). Yet the phase matters because it shapes how trajectories are explained, which policies can accelerate growth, and which dynamics can be extrapolated into the future. This raises the second question: what tools and metrics can identify phases, turning points, and the speed of diffusion without assuming a functional form in advance (RQ2)?

The third gap, in projection, is the widest. The dominant models say what a target requires or what would cost least, not what is likely. The data-driven alternatives have mostly extrapolated a single smooth trajectory (section 2.4). This raises the third question: what methods can probabilistically project the global growth of technologies from past data, and how do the resulting trajectories compare with climate-mitigation pathways (RQ3)? Chapter 3 sets out the framework, the data and the tools I use to take up these questions.

CHAPTER 3

Approach and methods

This chapter sets out the framework, data, tools and model developed in this thesis. The framework and the data come first (sections 3.1 and 3.2); Table 3.1 then summarises the methods, and sections 3.3–3.6 describe them.

3.1 Analytical framework: the four phases of technology growth

This thesis analyses technology growth through a framework that treats it as proceeding through distinct phases, rather than as a single smooth curve with a presupposed functional form. The classical S-curve tradition is retrospective: it studies technologies whose entire lifecycle – and therefore the final trajectory shape – is already observable. Instead, Paper I develops a framework that locates a technology, at a given point in time, in one of four phases (formative, accelerating growth, steady growth and slowdown), each shaped by a distinct configuration of drivers and barriers arising from the co-evolution of technology and policy (Figure 3.1). The framework was built for *policy-driven* technologies, whose deployment and growth are actively promoted by public policy to meet societal objectives such as climate change mitigation, energy security or industrial development. These contrast with *market-driven* technologies, which grow mainly through market forces with limited policy intervention (Paper I).

The first, formative phase is characterised by continuous experimentation, frequent failures, and substantial uncertainty. It ends with deployment ‘takeoff’, after which increasing economic and political returns progressively raise the speed of deployment in an accelerating growth phase. Acceleration ends at an inflection point, when

Table 3.1: Methods developed in this thesis, and the conventional practice each replaces.

Method	What it does
<i>Measuring growth</i>	
FOBI Section 3.3 · RQ2 Paper III	Detects takeoff from deployment data with piecewise Bayesian regression and model averaging. Instead of: fixed thresholds set as a deployment level or share of the eventual ceiling.
DAIP Section 3.3 · RQ2 Paper II	Detects the first inflection point and later growth pulses from annual additions, and measures G_{fp} and the cruising speed. Instead of: fitting a logistic curve, which averages over pulses and misdiagnoses the inflection point.
<i>Analysing costs</i>	
Phase-resolved cost analysis Section 3.4 · RQ1b Paper IV	Estimates learning and cost-reduction rates within each growth phase, using segmented regression mapped to phase transitions. Instead of: one rate fitted over the whole lifecycle.
<i>Comparing adopters</i>	
Three dimensions of diffusion Section 3.5 · RQ1c Paper V	Measures spread between countries separately from duration and magnitude within them, and relates each to lateness. Instead of: composite metrics that conflate the three.
<i>Projecting growth</i>	
PROLONG Section 3.6 · RQ3 Paper III	Projects global growth probabilistically from national growth parameters, combining simulated ensembles with quantile random forests. Validated by hindcasting against withheld data. Instead of: deterministic curve fits, scored on the same data they were fitted to.

emerging barriers such as grid integration challenges, supply-chain constraints, declining profitability and public opposition come to counterbalance these drivers. In the steady growth phase that follows, deployment continues at a quasi-linear rate, pulsing around a relatively stable cruising speed. Finally, in the slowdown phase, growth stalls as deployment approaches its ceiling: markets saturate, profitability falls as deployment moves to less favourable sites, innovation becomes incremental, newer technologies compete for market share,




Figure 3.1: The four phases of policy-driven technology growth. Figure 2(a) in Paper I.

and public and political support can fade (Paper I).

Paper I treats policy-drivenness as a binary attribute, although it more plausibly varies along a continuum and shifts as a technology matures. This thesis applies the framework to policy-driven technologies such as onshore wind and solar PV, and also tests how far it holds beyond them, for market-driven technologies such as mobile telephones and internet use. As Chapter 4 shows, the same phases appear in both, but policy-driven technologies reach the end of initial acceleration far earlier in their trajectories (Paper II).

3.2 Technology and country sample

This thesis analyses technology deployment at two levels: within individual countries, and in aggregate at the global scale.

At its core is an unbalanced panel of long-run deployment data for 91 technologies across 236 countries and dependent territories, extending from the nineteenth century to the present (Paper V). The series are compiled from statistics published by international organisations (the IEA, IRENA, the World Bank, FAOSTAT, UNICEF and the US Geological Survey), historical and academic datasets (CHAT, HATCH, the Banks cross-national timeseries archive and Mitchell's *International Historical Statistics*), industry databases (S&P Global's World Electric Power Plants database and the Energy Institute's *Statistical*

Review of World Energy), and other reports and grey literature¹. Each observation records a technology's deployment in a given country and year, in physical units such as installed capacity, vehicles or subscriptions, together with a matching market-size denominator such as total electricity generation or population. Deployment can therefore be analysed both in absolute terms and as a share of its market, and findings can be checked against the choice of measure. The raw compilation starts from 99 technologies in up to 243 countries; filtering retains 91 in 236. It keeps only country series that are at least 90% complete or have no more than five missing years, and only technologies for which at least five countries have an identifiable takeoff, so that countries can be compared within the same technology. Income and population from the Maddison Project, the World Bank and UN-DESA are added as covariates in Paper V. Papers II and III draw on subsets of this panel: electricity generation for wind, solar and nuclear power from the IEA, IRENA and Ember; gas turbine capacity from S&P Global; and mobile and internet use from the ITU and the World Bank.

The sample is deliberately broad. Alongside low-carbon energy technologies, it spans diverse sectors, eras and technology characteristics, from nuclear power and gas turbines to consumer technologies such as mobile phones and internet use. For Paper V, each technology is classified by sector, complexity and need for customisation [76], so the analysis can test whether a pattern holds across these characteristics. This diversity is what lets the analysis identify general regularities in technology growth, rather than the features of any single technology.

To study how technology growth co-evolves with cost, I use a second dataset, compiled for Paper IV, that links cost to cumulative deployment. Its series come from 64 peer-reviewed studies, institutional reports and public repositories. A technology is defined as a specific, repeatable artefact or process with a stable measurement boundary over time. This excludes broad sectors and fuel categories, whose costs can reflect input prices or product mix rather than learning. Of 213 candidate series covering 125 technologies, 145 met three criteria: cost and deployment refer to the same technology over the same years; the series begins at a low enough level of deployment to capture early or intermediate growth; and it contains at least five cost observations. The final sample covers 61 technologies, in 34 global and 111 national

¹Sources are named but not cited individually here. The complete list is available in the supplementary information to Paper V; Papers I–IV document their own subsets.

series, with deployment measured as cumulative production, installed capacity or adoption.

3.3 Model-agnostic tools for measuring technology growth

Situating a technology within the four phases requires identifying the transitions between them: the onset of accelerating growth, the shift to the steady growth phase and the oscillations within it. Typical curve fitting methods do not allow for this granular delineation of phases: they measure an asymptote approximating the deployment ceiling or ‘saturation level’, the inflection point marking a shift from accelerating to decelerating growth, and occasionally define takeoff at 2.5% of the final ceiling [58]. Given their reliance on knowing the ceiling, they require a timeseries covering the entire lifecycle and are thus unsuitable for technologies that are still growing. I therefore develop tools that can identify the shifts between phases directly from empirical data without presupposing a functional form.

The first, FOBI (Formative phase Bayesian changepoint Indicator), is built atop the Rbeast [160] algorithm and identifies the takeoff that signals the end of the formative phase and the beginning of sustained accelerating growth using a combination of piece-wise Bayesian regressions and model averaging to identify structural breaks in the timeseries where deployment switches from erratic fluctuations characteristic of the formative phase to consistent growth (Paper III).

The second, DAIP (Deceleration and Acceleration dynamics for Inflection point and Pulse detection), operates later in the lifecycle (Paper II). It identifies the first inflection point marking the shift from accelerating to steady growth where the annual rate of additions first peaks, and then locates subsequent pulses of deceleration and re-acceleration that punctuate the steady growth phase. I operationalise this in DAIP as follows. Starting from cumulative deployment D_t in year t , I compute annual additions ΔD_t , smooth them over a trailing window of ω years to give the deployment speed v_t , and take its first discrete derivative as the acceleration a_t :

$$\Delta D_t = D_t - D_{t-1}, \quad v_t = \frac{1}{\omega} \sum_{i=0}^{\omega-1} \Delta D_{t-i}, \quad a_t = v_t - v_{t-1}. \quad (3.1)$$

The first inflection point t_0 is the first year after takeoff where acceleration stays negative for n consecutive years, and a pulse begins at the first year t_p after t_0 where deployment is expanding and acceleration turns positive for n consecutive years:

$$t_0: \quad a_t < 0 \quad \text{for } t \in [t_0, t_0 + n - 1], \quad (3.2)$$

$$t_p: \quad v_t > 0 \text{ and } a_t > 0 \quad \text{for } t \in [t_p, t_p + n - 1]. \quad (3.3)$$

Paper II uses a three-year window ($\omega = 3$) and a two-year persistence criterion ($n = 2$); the persistence criterion prevents a single-year slowdown from being read as a phase transition, and Paper II tests the sensitivity of the results to both parameters. Because both criteria look backwards, DAIP diagnoses a transition only once it has occurred.

These tools yield a set of model-agnostic metrics that characterise a technology's growth: the timing of takeoff, the timing of subsequent growth pulses, and two rates. The first is the deployment speed at the first inflection point, G_{fip} , which is the maximum rate of annual additions reached during the initial acceleration phase. The second is the *cruising speed* $\bar{v}$, the mean annual additions from the first inflection point to the onset of sustained slowdown t_s , which Papers I–III use to characterise the steady growth phase:

$$G_{\text{fip}} = v_{t_0}, \quad \bar{v} = \frac{1}{t_s - t_0} \sum_{t=t_0+1}^{t_s} \Delta D_t. \quad (3.4)$$

Both rates are expressed in physical units per year or, when deployment is normalised, in percentage points of the relevant market per year.

3.4 Quantifying cost and learning dynamics

There have long been parallel traditions studying the dynamics of technology diffusion and the evolution of technology costs. This thesis brings the two together and asks how cost dynamics unfold across the phases of growth. Paper IV estimates two complementary measures of technology progress, both from log-transformed regressions within each phase. The learning rate follows Wright's law, relating unit cost C_t to cumulative deployment Q_t , and the annual cost-reduction rate follows Moore's law, relating cost to time [69]:

$$\log C_t = \alpha - w \log Q_t + \varepsilon_t, \quad \text{LR} = (1 - 2^{-w}) \times 100, \quad (3.5)$$

$$\log C_t = \alpha - m t + \varepsilon_t, \quad \text{CRR} = (1 - e^{-m}) \times 100, \quad (3.6)$$

where w is the learning exponent, m the annual cost-decline exponent, LR the percentage cost-reduction per doubling of cumulative deployment and CRR the percentage reduction per year. The two are linked through the growth rate of cumulative deployment, $g = d \log Q / dt$:

$$-\frac{d \log C}{dt} = w \frac{d \log Q}{dt} \implies m = w g. \quad (3.7)$$

The two rates therefore track one another only while g is constant, that is, while cumulative deployment grows exponentially². Once growth decelerates, a technology can keep learning at an undiminished or rising rate per doubling while its costs fall more slowly each year.

Rather than following the conventional practice of fitting a single rate over the full technology lifecycle, Paper IV estimates these rates phase by phase. It uses segmented regression to detect structural breaks in cost trajectories and map them to growth phase transitions identified using FOBI and DAIP. This makes it possible to ask whether learning and cost-reduction rates stay stable over the technology lifecycle, as techno-economic models and analyses typically assume, or shift as a technology matures.

Paper IV then examines how these patterns vary with the characteristics of a technology, particularly its complexity and its need for customisation [76]. Finally, because cumulative deployment and elapsed time tend to rise together, learning-by-doing is difficult to separate from progress that would occur over time regardless of deployment. A learning rate that rises over the lifecycle could reflect genuine learning, or simply time-driven cost decline measured against slowing deployment. To separate the two, Paper IV generates 10,000 synthetic technologies, each following a noise-perturbed S-shaped deployment trajectory, and simulates their costs under known mixes of the two mechanisms. The mixes run from purely time-driven decline to purely deployment-driven learning with an intensifying learning rate, plus a control in which the learning rate stays constant. To test which mechanisms reproduce the observed patterns, the same phase detection and rate estimation are then applied to the synthetic technologies as to the empirical data.

²This relationship between Wright's and Moore's laws was first described by Sahal [161].

3.5 Analysing diffusion between and within countries

Technologies do not spread uniformly across the world, and the commonly used terminology around “diffusion” tends to bundle together several distinct processes. I use *diffusion* as an umbrella term and separate three dimensions of it, which differ in their unit of analysis. I measure each with its own metrics, using deployment both in absolute units and normalised to the size of the relevant market. The first dimension, *spread*, is a process between countries: how quickly a technology reaches new adopters. It yields one value per technology. I identify the year each country takes off using FOBI, measure spread as the 10th-to-90th-percentile range of country takeoff years t^{off} ,

$$S_j = t_j^{\text{off},(90)} - t_j^{\text{off},(10)}, \quad (3.8)$$

and regress that span on the year the technology first took off anywhere.

The other two dimensions are measured within each country and then compared across countries. The second, *duration*, is the length of a technology’s passage through the lifecycle within a country. Using DAIP to locate its first and last inflection points, I measure the length of the formative, initial acceleration and steady growth phases, alongside the conventional transition duration Δt from 10% to 90% of a fitted logistic ceiling, and the growth constant k of that fit. The third, *magnitude*, is how fast and how far deployment goes within a country. I measure it through the compound annual growth rate during initial acceleration, the peak growth rate at the first inflection point, the cruising speed through steady growth, and peak deployment.

To compare duration and magnitude between early and late adopters, I measure each country’s lateness as the number of years between its takeoff and that of the technology’s first adopter. I then regress each within-country outcome Y_{ij} , for country i and technology j , on its relative timing T_{ij}^{rel} , the number of years between that country’s takeoff and the pioneer’s, absorbing technology fixed effects γ_j :

$$Y_{ij} = \beta T_{ij}^{\text{rel}} + \gamma_j + \varepsilon_{ij}, \quad T_{ij}^{\text{rel}} = t_{ij}^{\text{off}} - \min_{i'} t_{i'j}^{\text{off}}. \quad (3.9)$$

The fixed effects restrict identification to the between-country, within-technology relationship, so that later adopters are compared only with

earlier adopters of the same technology, and β is the change in outcome per additional year of lateness. Rates and levels are analysed in logs, so their coefficients read as semi-elasticities, while phase durations, the initial acceleration CAGR, k and Δt are analysed in levels. Standard errors are clustered two ways, by country and by technology, across 7,130 country–technology observations. I report three specifications: pooled OLS without fixed effects, the fixed-effects specification above, and that specification with log GDP per capita at takeoff added, to test whether the timing effect is a proxy for income.

3.6 Projecting technology growth with PROLONG

The metrics and phase diagnostics developed so far describe how technologies have grown in the past. To ask how they are likely to grow in the future, I develop PROLONG (PRObabilistic model Of techNo-logy Growth), a data-driven model that links the empirically observed growth patterns of individual countries to eventual global outcomes (Paper III). The relationship between early national trends and later global diffusion is hard to establish directly from historical data, especially for technologies that are still in the early phases of growth. PROLONG addresses this by combining Monte Carlo simulation with machine learning to explore that relationship systematically. It formalises three of the regularities this thesis identifies: growth proceeds through phases shaped by different mechanisms, so observations must be segmented by phase before they can inform extrapolation; technologies often show prolonged steady growth composed of several pulses; and the mechanisms of growth are broadly similar across early and late adopters, so front-runners can inform expectations for the rest.

The model is built and applied in five steps (Figure 3.2). First, for each technology, I simulate a large ensemble of 13,000 “virtual worlds”, each composed of 150 countries which correspond to real-world sizes. The technology takes off in every country in a different year and follows a noise-perturbed logistic or bilogistic [162] trajectory:

$$y(t) = \frac{L}{1 + e^{-k(t-t_0)}}, \quad G = \frac{kL}{4}, \quad \Delta t = \frac{\ln 81}{k}, \quad (3.10)$$

$$y(t) = \frac{L_1}{1 + e^{-k_1(t-t_1)}} + \frac{L_2}{1 + e^{-k_2(t-t_2)}}, \quad (3.11)$$




Figure 3.2: Five-step approach for PROLONG. a, Simulation and model training (Steps 1–3). b, Projection and validation with empirical data (Steps 4–5). Figure 3 in Paper III.

where L is the growth asymptote or deployment ceiling, k the growth constant, t_0 the inflection point, G the maximum annual growth rate and Δt the conventional duration of transition from 10% to 90% of the ceiling. The growth parameters are drawn from broad distributions, which let the ensembles span not only historically observed dynamics but also as yet unseen but plausible ones. The bilogistic worlds split growth into two pulses, so that it slows and rebounds to produce a prolonged steady growth phase that a single S-curve cannot represent. Second, I generate training data by truncating the simulated national series at successive points and estimating two growth parameters from each. FOBI first identifies takeoff, and the formative phase is excluded because it carries no information about later growth. A logistic function is then fitted to the post-takeoff data to recover the growth constant k and, for countries whose fitted curve has passed half its ceiling (the criterion by which Paper III identifies the inflection point), the peak growth rate $G = kL/4$. The logistic serves here only to read these two parameters. PROLONG does not use L because curve fitting cannot reliably estimate the ceiling from an incomplete series. It uses k and G because their distributions are stable under truncation. Paper II also shows that G overlaps strongly with the model-agnostic G_{fip} even where the two methods date the inflection point differently. Third, I train quantile random forests [163] on these data to learn how such early national signals map onto the eventual global growth constant and peak growth rate, using the median and interquartile range of k and G across mature countries, together with their number and their combined share of the global market as training features.

Applied to real-world data (step four), the same features are estimated from observed national series, and the trained forests predict global k and G probabilistically. Sampling national parameters across quantiles yields projection intervals rather than a single deterministic trajectory. PROLONG projects growth speed, not the ceiling: the ceiling it reports, $L = 4G/k$, is an indicative by-product and not an estimate. Finally (step five), I validate the model by hindcasting, splitting the timeseries and testing whether projections made from early data reproduce withheld later data, for onshore wind and solar PV as well as for reference technologies chosen to span diverse characteristics and diffusion timescales (nuclear power, combined-cycle gas turbines, mobile telephones). I benchmark it against global logistic and exponential extrapolations, against the aggregation of national

logistic projections, and against IEA and IPCC scenarios [3, 164, 165]. Variants trained on logistic-only, bilogistic-only and mixed worlds are compared the same way, and the best-performing variant is retained for each technology.

Beyond these probabilistic projections, I build counterfactual scenarios in which trajectories depart from PROLONG's median projection to reach, by 2040, the median of the IPCC AR6 1.5 °C pathways with no or limited overshoot [164] or the IEA Net Zero by 2050 scenario [165], and distribute the required growth across ten world regions under empirical constraints on ceilings, growth and acceleration. I use these to quantify how much faster growth would need to proceed in different global regions to meet stated climate goals.

CHAPTER 4

Results

4.1 Understanding how technology growth unfolds

4.1.1 The four phases of growth and the co-evolution of technology and policy (*RQ1; Paper I*)

Paper I proposes that the growth of policy-driven technologies proceeds through four distinct phases: formative, accelerating growth, steady growth, and slowdown. Each phase is produced by a distinct configuration of drivers and barriers arising from the co-evolution of technology and policy. The phases are not merely descriptive stages of a curve but the visible signature of an underlying interaction between what a technology can do and how policy responds to it.

Technology and policy co-evolve through both reinforcing and countervailing feedbacks. In the early phases, deployment and policy reinforce one another: deployment lowers costs and improves performance, which strengthens the case for support, which in turn catalyses further deployment. Deployment also creates constituencies: the takeoff of wind power in Germany in the 1990s built the interest groups that later lobbied for stronger renewables policy [89]. But the same expansion that sustains growth eventually generates resistance. As a technology becomes more pervasive, it begins to impinge on incumbents, landscapes, and existing systems, mobilising opposition from the actors it threatens or displaces, and this can weaken, stall, or reverse the support on which growth depends. Growth is therefore neither self-sustaining nor self-limiting in any simple sense: it is the moving balance of these forces, which shifts as a technology matures.

Policy support does not simply decline once a technology becomes

commercially viable; it often expands and diversifies as new barriers emerge. The barriers themselves change character across the phases: from the research, demonstration, and market-creation challenges of the formative phase, through the reinforcing feedbacks of acceleration, to the integration, social, and political barriers that surface once deployment reaches significant scale. Paper I traces this for onshore wind in Germany and solar PV in India. In both countries the number and scope of active policies grew as the technology matured, and the policy mix broadened from research and market creation to land use, acceptance, system integration, and domestic manufacturing, drawing in a wider set of ministries. Financial support per kilowatt-hour fell, but more slowly than the levelised cost of electricity. Total support fell more slowly still, because falling unit support was offset by a growing deployment base. The one break in the pattern came in Germany between 2022 and 2024, when the energy crisis pushed wholesale prices above guaranteed tariffs and public support briefly collapsed.

Because the phases differ in their dynamics, acceleration means something different in each. In the formative phase it means earlier takeoff. In accelerating growth it means a longer or steeper period of quasi-exponential growth. In steady growth it means a faster cruising speed, or the same speed sustained for longer. In slowdown it means restarting growth. Paper I also singles out the steady growth phase as the one in which technologies reach their fastest sustained growth and in which most of their eventual deployment is decided.

4.1.2 Growth pulses and the departure from smooth S-curves (*RQ1a; Papers II, III*)

If growth proceeds through phases, the next question is what those phases look like in the data. Applied to national deployment trajectories, the diagnostic tools described in section 3.3 show that growth departs from the smooth S-curve in two specific ways. Initially accelerating growth ends far earlier than classic S-curves like the logistic model imply, and what follows is not a monotonic slowdown but a prolonged phase of steady growth punctuated by repeated decelerations and re-accelerations (Figure 4.1).

The first departure concerns *when* acceleration ends. Takeoff, the end of the formative phase, occurs at around 0.8% of national electri-




Figure 4.1: Growth pulses lead deployment to deviate from smooth S-curves. Adapted from Figures 1 & 2 in Paper II.

city generation for onshore wind (IQR 0.5–1.4%) and 0.7% for solar PV (IQR 0.3–2.1%). Globally, both took off when they first exceeded 0.15% of generation, wind in 1999 and solar in 2010. The first inflection point, the year in which annual additions first peak, arrives about six years after takeoff for wind (IQR 3–9 years) and three years for solar (IQR 2–5), when the technology has reached roughly 3% of generation (IQR 2–8% for wind, 1–7% for solar). A logistic curve places its inflection at exactly half the eventual ceiling, so an inflection at 3% would imply that wind and solar were destined to saturate at around 6%. For ceilings of 30–50% of generation, the logistic model would instead expect the inflection at 15–25%. Deployment already comfortably exceeds 6% in many countries, and Denmark generates more than half its electricity from wind. The first inflection point therefore does not mark the midpoint of anything. It marks the end of initial acceleration.

The second departure concerns what happens afterwards. Under a standard S-curve, passing the inflection begins a monotonic decline in

annual additions that continues until deployment hits its ceiling. The data show something else. Deployment settles into a prolonged steady growth phase in which annual additions continue at a quasi-linear rate, the cruising speed. That speed is about 0.7 percentage points of generation per year (IQR 0.4–1.5 for wind, 0.3–1.9 for solar), slightly below the rate reached at the first inflection point. The steady growth phase is punctuated by recurrent *pulses*: episodes of deceleration followed by renewed acceleration. The growth rates in secondary pulses sometimes peak above those in the first, but mainly in countries whose first peak was slow, so the average post-inflection rate stays below the first peak in most countries. Italy's solar boom of 2007–2012, its collapse after retroactive tariff cuts, and its restart after 2020 is one such pulse; Denmark's onshore wind stagnation after the feed-in tariff was cancelled around 2000 and its re-acceleration after support returned in 2008 is another (Paper I). Growth after the first inflection point is better described as sustained and episodic than as smoothly decaying.

The two departures are not equally generalisable. Steady growth around a cruising speed close to the first peak, with pulses, appears across nuclear power, combined-cycle gas turbines, mobile telephones, and internet use as well as wind and solar. However, mobile telephones reach their first inflection at a median of 79% of per-capita penetration and internet use at 38%, roughly an order of magnitude later in the trajectory than the 3% observed for renewables, and the post-inflection slowdowns of the older technologies are more pronounced. Paper II hypothesises that the difference reflects the larger influence of policy on renewables. Policy is more fluid and adaptive than the techno-economic drivers, and the same volatility that stalls growth through rollbacks or bottlenecks can restart it through renewed commitment, producing early inflections and later pulses rather than a single smooth curve.

4.1.3 Cost and learning dynamics across the phases (*RQ1b; Paper IV*)

Paper IV structures cost data by growth phase rather than calendar time. It assembles 145 cost and deployment series covering 61 technologies, 34 at the global level and 111 at the national level, which more than doubles the evidence base of earlier studies and, for the first

time, covers, for half the sample, all three phases before slowdown. Two rates are estimated within each phase: the learning rate, the percentage cost decline per doubling of cumulative deployment, and the cost-reduction rate, the percentage cost decline per year. The central finding is that the two diverge once technologies enter the steady growth phase (Figure 4.2). Global learning rates rise from a median of 15% (IQR 13–21%) in the formative phase to 26% (22–35%) in accelerating growth and 32% (24–43%) in steady growth. Cost-reduction rates rise from 10% (6–15%) to 14% (10–16%) and then fall to 3% (2–9%) once technologies enter steady growth. The national sample shows the same shape at lower levels: learning rates of 10%, 19%, and 23%, and cost-reduction rates of 5%, 7%, and 5%. Within individual trajectories, 71% of paired series show a rise in the learning rate from the formative to the accelerating phase and 59% from the accelerating to the steady growth phase, while cost-reduction rates fall across the second transition in 68% of series. The early rise in cost-reduction rates is the less consistent half of the pattern: rates increase across the first transition in only 49% of paired series, so it is the later slowdown that the evidence supports most strongly. The rise in learning rates between the formative and accelerating phases is significant in both samples; the rise into steady growth is significant nationally but not globally, where only five technologies have reached that phase.

The divergence is not obvious, because the two rates are often treated as interchangeable. They are equivalent only under sustained exponential growth, where the annual cost-reduction rate is simply the learning rate multiplied by the annual growth rate of cumulative deployment. Deployment does not grow exponentially for long. Year-on-year growth rates fall systematically across the phases: in the global sample, the median drops from 58% (IQR 31–112%) in the formative phase to 27% (16–53%) in accelerating growth and 8% (4–12%) in steady growth, and the national sample falls from 76%, to 28%, to 9%. Once growth decelerates, each further doubling of cumulative deployment requires larger absolute additions, and takes longer to achieve.

Learning and deployment growth therefore operate as competing feedback loops. A reinforcing loop drives costs down, as cheaper units spur further deployment and further experience. A balancing loop weakens that reinforcement as slowing deployment growth means experience accumulates more slowly against a larger base. Annual cost




Figure 4.2: The evolution of learning rates (a, b), cost-reduction rates (c, d), and year-on-year deployment growth rates (e, f) across the phases of growth. Figure 3 in Paper IV.

reductions follow a hump-shaped trajectory: high in earlier phases when deployment is growing fast, then declining even as learning per doubling continues to improve. The empirical pattern alone cannot say whether the rising learning rates reflect learning-by-doing or are an artefact of time-driven progress measured against slowing deployment. Paper IV separates the two with computational experiments on 10,000 synthetic technologies whose cost declines are generated under known mixes of the two mechanisms. Only the scenarios combining intensifying learning-by-doing with autonomous time-driven progress reproduce the full pattern. Purely time-driven costs produce a rising learning rate but a flat cost-reduction rate; purely learning-driven costs with a fixed learning rate produce a monotonically falling one.

Two further observations support a phase-based reading of these dynamics. First, structural breaks in learning and cost-reduction rates fall within a median of three to five years of a growth phase transition. That is significantly closer than a Monte Carlo null in which break-points are placed at random ($p < 0.05$ for all technology types). The one exception is national series for complex, customised technologies

such as nuclear power, where sparse deployment makes both transitions and breaks harder to locate. Second, the reason this variation went undetected is a data constraint, not its absence. Most earlier analyses rested on series that capture only the formative phase, before any systematic pattern would be expected to emerge. Paper IV's larger sample also confirms, now with statistical significance, that learning rates fall with design complexity, from 21% for standardised mass-produced technologies to 15% for platform-based and 6% for complex customised ones.

4.1.4 Early and late adopters: spread, duration, magnitude and absorptive capacity (*RQ1c; Paper V*)

Section 2.2 described a long-standing debate on whether latecomers adopting new technologies gain from the proven designs and lower costs they inherit or lose from their weaker capacity to take them up. The argument remains unresolved because the standard metrics used to measure growth lump together different things. Paper V separates them into the three dimensions distinguished in section 3.5: the spread of a technology between countries, and the duration and magnitude of its growth within each one. Cross-country spread is measured across technologies. Duration and magnitude are measured for 7,130 country–technology pairs, covering 91 technologies in 236 countries and territories over two centuries, and compared between early and late adopters. Analysed separately, the three dimensions move in different directions. Technologies reach new countries ever sooner, and later adopters pass through shorter lifecycles, but they deploy more slowly and saturate at lower levels (Figure 4.3). Each side of the old argument turns out to have been right about a different dimension.

Cross-country spread is narrowing: technologies reach most adopting countries within an ever shorter span of years. The span of years within which the middle 80% of adopting countries take off has shortened by 5.5 years for every decade later a technology first appeared ($n = 52$ technologies, $R^2 = 0.87$). The telegraph took more than a century to take off in most countries; mobile telephones and solar PV managed it in little over a decade. Paper III measures the same thing for solar PV and onshore wind directly. The time for the number of countries past takeoff to grow from 10% to 90% of its estimated




Figure 4.3: Technologies reach new countries ever sooner, and later adopters pass through shorter lifecycles, but they deploy more slowly and saturate at lower levels. Adapted from Figure 1 in Paper V.

maximum is 23 years for onshore wind and 15 for solar, faster than for any other electricity technology in its sample.

Within countries, later adopters compress the technology lifecycle. Each additional year of lateness is associated with a shorter formative phase ($\beta \approx -0.30$), a shorter acceleration phase (-0.15), and, most strongly and most consistently, a shorter steady growth phase (-0.55). Transition durations (the duration it takes deployment to grow from 10% to 90% of the final ceiling) shorten accordingly (-0.26 years per year of lateness), and because duration and steepness are inversely related, the fitted logistic growth constants of later adopters are correspondingly higher ($+0.004$). This compression is behind the optimistic reading, in which countries that adopt technologies later

move through the lifecycle faster [34, 57].

The pattern reverses along the dimension of magnitude. Initially accelerating growth shows no robust relationship with lateness. But each year of lateness is associated with a roughly 0.8% lower peak growth rate at the first inflection point and a 1.4% slower cruising speed through steady growth. These penalties compound over time: a country taking off twenty years after the pioneer grows about 25% more slowly over the phase in which most deployment occurs. More consequentially, later adopters saturate lower. Peak deployment falls by about 2% per year of lateness, so a country with a twenty-year lag peaks at roughly two-thirds the deployment share of an early adopter. These associations are essentially unchanged when income at takeoff is controlled for, so the penalty does not simply reflect differences in country affluence, and they hold across the complexity spectrum, from vaccines and radios to nuclear plants and steelworks. The two findings together resolve the contradiction between late-comers' higher fitted growth constants and slower actual deployment: the growth constant measures the steepness of a curve that is both compressed in time *and* capped at a lower ceiling.

The data suggest that this ceiling is more strongly influenced by countries' absorptive capacities than by the timing of adoption. The clearest evidence comes from mobile telephones, the exemplar of leapfrogging. Across 163 countries, peak mobile penetration exceeds 100% almost everywhere, but it rises with prior landline penetration ($r = 0.32$), and with income at takeoff, and the gradient is steepest within Africa. Each predictor is significant on its own. Regressed together, income remains significant while prior landlines and timing do not. The pooled regressions show that the lateness penalty survives controlling for income; the mobile case shows that, for ceilings, income and prior infrastructure carry more of the explanation than timing. Read together they say that what looks like a penalty for lateness is largely a penalty for low absorptive capacity. Leapfrogging is real but relative. A country can leap over its own past, but not above others with greater capacity to deploy and integrate a new technology.

These patterns are not confined to the past. Solar PV has taken off in scores of countries, with most adopters taking off within a decade of each other. Later adopters already finish initial acceleration sooner, by around four years per decade of delay relative to the pioneer, without any significant change in growth constant, acceleration phase growth

rate, or peak growth rate. Electric vehicles show the same temporal compression: the time for electric cars to go from 1% to 10% of new sales fell from 5.2 years for countries crossing 1% before 2015, to 4.3 years for the 2015–2018 cohort, to 3.5 years from 2019 onwards. Because vehicle fleets turn over slowly, sales run far ahead of stock, at 97% of new-car sales in Norway against 36% of its fleet, and meaningful EV adoption is observed in only about 60 countries, presumably those with the highest absorptive capacity. It is too early to observe where either technology will saturate. If history is a guide, falling costs will take them almost everywhere, and countries' capacities will decide how fast and how far they go.

4.2 Measuring the dynamics of technology growth

4.2.1 Detecting takeoff, inflection points, and growth pulses (*RQ2; Papers II, III*)

The phases described in section 4.1 are only useful if they can be identified for technologies whose growth is still under way. Applied across the technology sample, the diagnostics of section 3.3 identify each transition directly from deployment data, without an assumed functional form and without knowledge of the eventual ceiling.

For takeoff, FOBI identifies country-specific estimates of when sustained growth begins. FOBI thereby departs from the convention of defining takeoff by a fixed threshold, such as a given share of the eventual market, which applies one criterion to countries of very different sizes and starting positions (Paper III). Applied to 2023 data, FOBI diagnoses takeoff in 63 countries for onshore wind and 90 for solar PV.

For the phases that follow, DAIP identifies the first inflection point and the subsequent pulses of deceleration and re-acceleration, and does so more reliably than curve fitting. Paper II systematically compares DAIP and curve fitting using a confusion matrix. It finds that curve fitting both reports “phantom” inflections in countries that are still accelerating and misses real ones in countries whose inflection is followed by a pulse. The reason is structural. A single logistic fitted to a pulsating series averages across the pulses, and so either misses the inflection or misplaces it. Fitting logistic curves to progressively truncated series demonstrates how the fitted inflection point shifts forward whenever a secondary pulse occurs. Curves with several in-

flection points do not solve the problem, because they require the number of pulses to be assumed in advance, and DAIP diagnoses a transition only once two years of data confirm it. Smooth curves are not merely imprecise about the timing of pulses, but are structurally incompatible with them.

Not all growth metrics are equally reliable. The deployment ceiling estimated from an incomplete series is highly unstable and shifts substantially as the series is truncated at different points, whereas the peak growth rate is comparatively robust and can be inferred reliably from early national experience (Paper III). The growth rate at the first inflection point (G_{fip}) approximates the maximum growth rate G of classical growth models while being less sensitive to pulses, which makes it the more dependable of the two (Paper II). Year-on-year growth rates, and metrics resting on exogenously assumed asymptotes or transition durations, prove poor descriptors of growth in this setting: the first because they are sensitive to short-run volatility, the others because they are sensitive to the exogenous parameter assumptions. Projecting constant annual additions at G_{fip} tracked subsequent deployment more closely than exponential extrapolation, than logistic curves with empirically estimated ceilings, which systematically underpredict growth, and than logistic curves with ceilings fixed at policy targets, which perform almost as poorly as exponential fits. The results are robust to reasonable variation in DAIP's smoothing window and persistence criterion, although the number of pulses detected is more sensitive than the distributions of G_{fip} and cruising speed. The trustworthiness of a metric has direct implications for how reliably it can inform projections.

4.2.2 Cruising speeds and the relationship between national and global growth (*RQ2; Paper II*)

Most deployment happens in the steady growth phase, which makes the cruising speed especially salient for anticipating how far a technology will grow. Section 4.1.2 showed that cruising speed sits close to, but slightly below, the growth rate at the first inflection point. What matters for measurement is that G_{fip} can be estimated robustly, and long before the steady growth phase ends, so it serves as a reliable predictor for the speed a country will sustain through that phase.

Countries that have already entered steady growth therefore consti-




Figure 4.4: The relationship between national G_{nip} and global growth rates. Adapted from Figure 5 in Paper II.

tute a sample of observed feasible cruising speeds. The distribution of their peak rates is stable over time, and it defines an upper bound for global deployment speed long before global deployment reaches its own inflection point. The sample becomes reliable once two conditions are met: at least five countries have passed their first inflection point, and the technology has taken off globally.

Wind and solar in later-adopting countries rarely grow faster than in front-runners, despite the advantages latecomers enjoy in technology cost, supply-chain maturity, and accumulated experience elsewhere. Those advantages appear to be offset by less favourable geographic, infrastructural, economic, institutional, and political conditions. The net result is that the upper bound of national growth rates has not risen.

Taken together, these results link national and global growth. Global deployment is the aggregate of national trajectories, so if the fastest national cruising speeds stay stable, the median and upper quartile of front-runner peak rates form an empirically grounded frontier for plausible global growth (Figure 4.4). For onshore wind, nuclear power, gas turbines, internet use, and mobile telephones, the median national G_{nip} has consistently exceeded the observed global growth rate once enough countries had passed their inflection, and the up-

per quartile has provided a more optimistic ceiling. Solar PV is the exception, for now. Global solar growth currently tracks the upper quartile of the national distribution rather than its median, because China accounts for the majority of global additions while growing faster than the cross-country median on the strength of its manufacturing expansion. As growth in China and other large accelerating economies stabilises, Paper II expects the solar relationship to settle where the wind relationship already has. This has two implications. Sustained exponential growth at the global level is inconsistent with the evidence, in line with recent near-linear projections from the IEA and BloombergNEF [3, 166]. And because a few large countries can dominate global totals, moving from national distributions to global aggregates requires weighting countries by their share of the global system. Paper II also compares the national G_{fix} distribution against the maximum annual additions in IPCC AR6 1.5 °C and 2 °C pathways, a comparison that section 4.3.2 develops with PROLONG.

4.3 Projecting data-driven trajectories for technology growth

4.3.1 Probabilistic projections of global wind and solar growth (*RQ3; Paper III*)

Paper III develops PROLONG (section 3.6) and applies it to onshore wind and solar PV. Applied to observed national data through 2023, PROLONG projects onshore wind at 13.4% of global electricity generation in 2030 (IQR 12.6–14.3%), 22.8% in 2040 (18.3–26%), and 25.6% by 2050 (20.5–33.7%), and solar PV at 12.3% in 2030 (10.6–14.8%), 19.6% in 2040 (14–27.5%), and 20.8% by 2050 (14.5–29%). At the 95th percentile, wind could reach 17%, 34%, and 47% and solar 18%, 35%, and 47% at the same dates (Figure 4.5).

The projections are validated by hindcasting. Paper III withholds later observations and projects forward from the data available at a given point in time, then compares the projection with what happened. It does this at horizons up to thirty years, for wind and solar and for three reference technologies chosen to span the spectrum of complexity and policy dependence: nuclear power, combined-cycle gas turbines, and mobile telephones. Across all five, PROLONG outperforms global logistic and exponential curve fits and the aggregation of

national logistic extrapolations. Its median projections keep a symmetric mean absolute percentage error typically below 0.5 and a symmetric mean percentage error close to zero across the full range of horizons, which indicates accuracy without systematic bias. Global exponential and logistic fits often exceed 1.0 on the first measure and show biases that worsen with horizon.

For instance, projecting from data through 2015, PROLONG's median put onshore wind at 7.8% of global electricity in 2023 (IQR 7.1–8.8%) and solar at 4.8% (3.6–6.3%). The actual outcome was 7% and 5.3%. Over the same horizon the exponential fit gave 13% for wind and 12% for solar, and the logistic fit 5% and 1.5%: one far too optimistic, the other far too conservative.

Paper III also assesses the calibration and sharpness of the intervals, through coverage rates, interval scores, the continuous ranked probability score, and diagnostics of tail behaviour. Out-of-sample outcomes fall inside the predicted ranges at close to the nominal frequencies. The intervals widen with horizon, as they should, but stay well behaved rather than expanding without limit, and the distributions avoid the heavy tails that would make them uninformative.

Two limits define how the projections should be interpreted. They are projections of growth rates rather than of saturation levels, so they describe how fast deployment is likely to proceed rather than where it will stop (section 5.4). And they are conditional on the future resembling the past in structure, if not in detail. They extrapolate the mechanisms visible in national experience, and cannot anticipate discontinuities that break with it.

4.3.2 A baseline for benchmarking climate-mitigation pathways (*RQ3; Paper III*)

By projecting what is likely given current trends, PROLONG offers a benchmark for climate mitigation pathways and policy pledges that describe what is required to meet the climate goals.

In hindcasts made in each year from 2008 to 2022, PROLONG's median projections were consistently higher, and more accurate, than the IEA's contemporaneous stated- and new policies scenarios and the baseline scenarios of the IPCC's Fifth Assessment Report and Special Report on 1.5 °C. For its latest projections starting in 2023, the IEA's stated policies scenario [165] closely tracks PROLONG's median for




Figure 4.5: Probabilistic projections and mitigation-aligned scenarios for global onshore wind (a, c) and solar PV (b, d) deployment. Black dots show historical data. Coloured point-range lines depict deployment distributions for current policy, baseline scenarios (a and b) and climate mitigation pathways (c and d) from the IPCC AR6 [164]. Dot–dash lines in (c) and (d) show early and late acceleration scenarios aligned with reaching deployment levels consistent with the IEA NZE scenario [165] by 2040. Figure 5 in Paper III.

onshore wind, while most baseline and current-policy scenarios in the AR6 database [164] fall below its 25th percentile. For solar, the AR6 baselines again fall below the 25th percentile and its current policy scenarios start below the median before converging by 2040 and the IEA’s WEO 2025 stated policies scenario [3] falls inside the interquartile range.

PROLONG’s interquartile range largely overlaps the 2 °C-compatible pathways of AR6 for both technologies, and the 1.5 °C high-overshoot pathways for onshore wind. The more ambitious 1.5 °C no- or limited-overshoot pathways, and the IEA’s Net Zero scenario, exceed PROLONG’s 95th percentile during 2030–2035, particularly for solar. The gap narrows after 2035, as the pathways approach saturation while PROLONG’s upper range continues to expand: by 2050 its 75th percentile encompasses even the most ambitious wind pathways and

approaches the median of the stringent solar ones. In the median projection, renewables grow more slowly, but over a longer period, than many 1.5 °C pathways require. Paper III also compares PROLONG to some of the more ambitious projections and scenarios in the recent literature [5, 6, 165, 167, 168]. While the comparison is mixed for onshore wind, for solar PV, all five of these lie above PROLONG's 95th percentile through 2040. This places PROLONG between baseline scenarios that have run consistently too low, and ambitious, cost-driven projections that now run above anything countries have sustained. The COP28 pledge to triple renewable capacity by 2030 sits near the 95th percentile of PROLONG's distribution. Reaching it would require the growth of onshore wind in major economies to accelerate by 1.4 to 3 times, and of solar PV by 2 to 5 times, relative to recent rates.

Paper III constructs two acceleration scenarios against a baseline aligned with PROLONG's median. Early acceleration departs upward immediately, meeting the COP28 pledge by 2030 and, by 2040, the median of the AR6 1.5 °C no- or limited-overshoot pathways or the IEA's Net Zero scenario. Late acceleration follows the baseline until 2030 and must then rise more steeply to reach the same 2040 level. Paper III distributes the required growth across ten world regions under empirical constraints on ceilings, growth, and acceleration. For onshore wind, early acceleration requires additions of about 1.8 percentage points of generation per year in Europe and East Asia and up to 2.7 elsewhere, rates similar to the European Union's REPowerEU plan. East Asia and North America would have to accelerate to 1.4 to 1.6 times their previous peaks, and other regions to between 2 and 14 times. For solar, every major region would need to exceed 2.7 percentage points per year, two to four times its historical peak and above China's record additions of 2024. Such rates are not unprecedented: India's Central Electricity Authority targets about 2.4 percentage points per year through 2030, up from 1 today. But they have been reached only in exceptional episodes, and the scenarios require them to be sustained over multiple years.

Discussion

5.1 Main contributions

Across five papers, this thesis shows that there is more to the growth of new technologies than how much they cost. Each paper challenges a different part of commonly held assumptions about how technologies grow: growth co-evolves with policy and proceeds through phases and pulses rather than smoothly (Papers I–III); costs do not fall at a constant rate but decelerate even as learning intensifies (Paper IV); and technology does not spread evenly but is shaped by the national capacity to absorb it (Paper V). Together they replace a smooth, cost-driven picture of growth with an empirically grounded, phase-based one that accounts for both techno-economic and socio-political mechanisms.

The conceptual contribution is a phase-based framework for analysing technology growth. Classical S-curves assume that acceleration is followed immediately by slowdown, which can overestimate a technology's prospects by extending the accelerating phase too far, or underestimate them by mistaking a temporary deceleration for saturation. This thesis proposes instead that the end of accelerating growth is often followed by a prolonged steady growth phase, in which deployment proceeds at a relatively stable cruising speed punctuated by pulses of deceleration and re-acceleration (Papers I, II). Growth thus falls into four phases – formative, accelerating growth, steady growth and slowdown – which differ not only in their dynamics but in the mechanisms that shape them. A second conceptual contribution unpacks technology diffusion into three dimensions: the spread of a technology between countries, and the duration and magnitude (speed and eventual extent) of its growth within each one (Paper V).

The methodological contribution is a set of model-agnostic tools for measuring technology growth and a probabilistic model for projecting

it. FOBI and DAIP delineate growth phases directly from empirical data, detecting takeoff (Paper III) and the first inflection point and subsequent pulses (Paper II) without assuming that trajectories follow a fixed functional form. PROLONG turns early national growth signals into probabilistic, hindcast-validated projections of global growth (Paper III). It occupies a middle ground between energy and climate optimisation and simulation models on one side and bottom-up industry forecasts on the other, telling us not what is required, cost-optimal or planned but what is most likely.

The empirical contribution is the expanded evidence base and the regularities it reveals. The thesis draws on long-run deployment records for more than 90 technologies across over 200 countries, and adds 145 cost and deployment series covering 61 technologies, which more than doubles the evidence base of earlier studies and, for the first time, covers, for half the sample, all three phases before slowdown (Paper IV). Across dozens of technologies, these data show regularities that no single case could: a cruising speed close to, and typically bounded by, the peak growth rate at the first inflection point; learning rates that rise over the lifecycle even as cost-reduction rates slow; and later adopters that move through the lifecycle in less time but deploy more slowly (Chapter 4).

5.2 Implications for modelling, policy, and development

This thesis has implications for how technology growth is *modelled*, how it is *governed*, and how the transition interacts with economic development. The first two often co-evolve: models shape the targets that policy is designed to meet, and the impact of policies is then evaluated against modelled baselines [121, 122, 169]. Development shapes the broader landscape in which technology and policy interact.

Three findings from this thesis bear directly on how technology growth is represented in models. Because growth proceeds through distinct phases, models that impose smooth deployment, whether by cost optimisation or simple extrapolation, misrepresent technology trajectories (Papers II, III). Because learning rates rise over the lifecycle while cost-reduction rates slow, two common ways of representing cost dynamics carry a predictable bias (Paper IV). Models that endogenise learning at a fixed rate [78] can understate how much

is learned in later phases. Stochastic forecasts centred on a constant rate of annual decline [6, 77] can overstate long-run cost reductions, because their central trend extends early-phase rates into the steady growth phase. Models could explicitly link learning to the phase a technology is in, with FOBI and DAIP locating that phase directly from deployment as it unfolds. Coupling both the learning rate and the cost-reduction rate to the growth phase could correct the bias in fixed-rate formulations.

Probabilistic projections such as the one developed using PRO-LONG in this thesis open a further possibility. Such projections could be soft-linked to integrated assessment models (IAMs) and energy system optimisation models (ESOMs), as probabilistic priors on deployment rates or as a check on the plausibility of modelled trajectories. Doing so at scale would require comparable probabilistic projections for other technology types, which could build on those already developed for mature technologies [27, 28] and on related work on technologies in the formative phase [138, 139, 170, 171]. Section 5.5 outlines the research this would require.

For policy, the central message is that effort does not fall as a technology matures. Barriers evolve rather than disappear – from the technological and economic barriers of the early phases to the integration, social, and political barriers of later ones – and thus support must be sustained and adapted even after a technology reaches cost-competitiveness (Paper I). The impact of different instruments shifts accordingly. Demand-pull helps catalyse growth in the early phases [172], while sustaining cost declines and growth through the steady growth phase depends on managing the barriers that emerge as technologies expand and begin to reshape the socio-technical systems they are embedded in (Papers I, IV). The shifting balance of mechanisms also reframes what acceleration means at different phases of growth: inducing takeoff in the formative phase; extending initial acceleration; countering emergent barriers to sustain and reintroduce pulses during the steady growth phase; and restarting growth – by expanding or re-defining the market or stimulating new innovation – after slowdown and saturation (Paper I). The finding that the growth rate at the first inflection point is frequently predictive of the overall cruising speed may also suggest that there are country- and technology-specific limits to sustained policy-driven acceleration of deployment. This is a hypothesis that future work should systematic-

ally test.

The findings of this thesis also show that history can offer an empirically grounded reality-check on targets. The frontier defined by front-runner peak speeds shows that meeting the COP28 pledge and longer-term Net Zero targets would require historically exceptional effort (Papers II, III). The acceleration required to meet the most ambitious goals cannot rely on technological progress alone; it requires sustained policy effort.

Whether that effort ultimately succeeds depends not only on policy design or technology cost but on the capacity of countries to absorb and deploy the technologies that the transition requires. Newer technologies reach latecomers sooner, but how far they grow within a country is bounded by that country's absorptive capacity. Paper V's finding reframes a longstanding debate about how technology latecomers catch up. Against the optimistic expectation of *leapfrogging* – that countries can vault over stages of development by adopting mature, cheap technology – the evidence suggests leapfrogging is real but relative: a country can leap over its own past, but not above frontrunners with higher absorptive capacity. The finding aligns with a structuralist reading of development, in which catch-up has depended less on *access* to technology than on the accumulated capabilities – skills, institutions, finance – needed to *use* it [53, 55, 56, 173].

Paper V also raises questions about global equity. The history of fossil-fuelled socio-economic development has long suggested a tradeoff between development and decarbonisation [30, 174]. The arrival of cheap low-carbon energy appears to have dissolved that longstanding tension [4, 6]. But if the *magnitude* of deployment is bounded by a country's capacity, then the binding constraint on equitable decarbonisation is not technology cost or carbon space but capacity itself. Closing the gap between early and late adopters is therefore, in large part, a development problem, and investment in capacity is a precondition for deep decarbonisation. Such investment spans human capital, including health, education, skills, as well as institutions, domestic finance, and system-planning capacity. Such investment also pays a second dividend, since the capacities that enable mitigation are largely those that enable adaptation and help build resilience to a changing climate [175, 176].

That investment is, however, in historic retreat. Contracting international aid, stalling progress on human development, and consistent

shortfalls in climate finance jeopardise the ability of many countries to take full advantage of this technological moment [177–180]. And because the hardest-to-decarbonise sectors often require the most complex, capacity-intensive technologies, without deliberate investment in capabilities, the clean-energy transition risks entrenching the very divide between developed and developing economies it might have helped to bridge.

5.3 Epistemological reflections

The implications I outline above lean on the regularities I identify in this thesis. The five papers have produced a good deal of knowledge about how technologies grow. In this section I step back and ask what kind of knowledge it is and to what use it can be put responsibly. The argument here runs in five steps: what kind of knowledge the thesis produces; how that knowledge could be used; what turning its patterns into explanations would require; what the vocabulary of diffusion assumes about them; and what the models built on that vocabulary, including this thesis' own, can and cannot see.

The thesis abstracts from deployment and cost data for many technologies in many countries to find regularities in how they evolve. In Sayer's terms the study design is extensive rather than intensive: it finds patterns common to a large population of cases, and trades the concrete causal history of any single case for that breadth [181]. Such a design does not yield laws. Following the critical realist tradition, I read the patterns here as demi-regularities: partial, approximate and context-bound tendencies of open socio-technical systems shaped by human agency [112, 182]. They recur often enough, and in enough contexts, to be informative, but they hold only while the conditions that produce them persist. The thesis identifies patterns and proposes explanations for them. It does not quantify causality.

Chapter 2 highlights the socio-technical tradition's emphasis on explaining why transitions happen (section 2.3). But even within the tradition, scholars disagree about what an explanation requires. Some ask for tested mechanisms and comparative designs [112]; others defend mid-range frameworks and processual accounts built case by case [108].

At first glance, this thesis appears to circumvent these debates by siding with the broader diffusion literature and focusing on measur-

ing growth. But there are two features of this work which also lean towards explanation. First, a robust regularity already signals that some mechanism is at work, even when it cannot say which. Second, the regularities this thesis finds are temporal. Phases, turning points and pulses describe how growth unfolds over time, the dimension which variable-oriented statistics capture least well, and the one socio-technical scholars place at the heart of explanation [108]. The thesis measures that temporal structure across more cases than any single process-tracing study could follow, but needs qualitative methods to explain it.

Knowledge of this kind lends itself to three uses. The first, most common use is to translate a tendency into a parameter and build it into a model. Wright's law – a regularity between technology costs and cumulative deployment – is a good example for how that happens. Observed often enough across different technologies, experience curves are now widely used to endogenise technology cost dynamics in integrated assessment and energy system models [78]. Historical deployment rates have travelled the same route and are often used to cap how fast new technologies may expand in scenarios [183, 184]. Some findings in this thesis, such as the phase-dependent learning and cost-reduction rates, could take the same route. However, efforts at assimilation also risk insulating a tendency from the conditions and contexts under which it holds, and the places where it breaks. Paper IV shows how commonly assumed constant, technology-specific learning rates are inconsistent with empirically observed patterns.

A second use is as an empirically grounded benchmark for what is plausible. A distribution of outcomes across comparable cases can serve this purpose without needing an explanation of causality. Section 5.2 used the frontier of front-runner growth this way, as a reality-check on the COP28 pledge and the most stringent mitigation pathways, and PROLONG's probabilistic projections rest on the same logic (Paper III). A benchmark inherits the caveat attached to every demi-regularity: it holds only while the conditions that produce the data persist (section 2.4). I return to what that caveat means for PROLONG at the end of this section.

The third use refers to what critical realists call retrodution: reading the gap between what existing theory predicts and what the data show as a clue to a mechanism the theory does not contain [112, 182]. For example, the S-curve's geometry implies there are only two phases,

acceleration and slowdown, with growth slowing as a technology nears its ceiling. This thesis finds growth slowing far earlier, and then pulsing between deceleration and re-acceleration rather than declining monotonically (Papers II, III). On the one hand, this observation could be used to specify a different functional form that yields a better-fitting curve. But on the other, it could be read retroductively and used to ask why this deviation occurs, and what begins to act on wind and solar at around 3% of electricity generation that was not acting before (Papers I, II).

Answering those questions means identifying mechanisms. This thesis proposes that the phases and their growth dynamics emerge from combinations of mechanisms whose causal powers may or may not activate in a given context. Public opposition is one such mechanism. It acts as a barrier to growth, and it has activated for onshore wind in parts of Northern Europe but has not for mobile phones. However, given that these processes play out over time, naming the causal mechanisms that can act on growth is not enough [112]; one must also trace how those actions unfold and interact over time [108].

Positing a causal mechanism is not the same as establishing it, and here the *extensive* design [181] of this work reaches its limit. The design can show where observations depart from expectation, but not which candidate mechanism produced a departure in a given country, nor whether that mechanism was necessary. Because mechanisms activate in some places and not others, the same phase-specific mechanisms yield different outcomes across countries, and discontinuities resist prediction. This is not for lack of data, but as a matter of ontology [112, 185]. Instead, through its extensive design, the thesis locates turning points, the onset of different phases across many countries, and so tells intensive, case-based research where and when to look. It also supplies the comparisons such research needs to judge whether a mechanism found in one case recurs in others (section 5.5). It does not dissolve the division between diffusion and socio-technical traditions described in Chapter 2 but facilitates an exchange in which each side poses questions the other can help answer.

Mechanisms are not autonomous. There is agency involved in activating them. Residents object, governments adjust support, firms invest. The language used to describe technology growth must therefore leave room for agency. Section 2.2 discussed the origins of the word *diffusion* and how it conjures images of a substance spreading

passively down a gradient, a contagion spreading through a population, or cultural traits radiating from a centre to a periphery. When used to describe the spread of technology, it smuggles in three assumptions about the nature of growth: passivity (a technology spreads by its own inherent properties), inevitability (an S-curve once begun always runs all the way), and a fixed direction of travel (what happens in central frontrunners will in time happen at the periphery).

The findings in this thesis contradict all three. Growth is shaped by a shifting balance of mechanisms rather than autonomous (Paper I), pulsating rather than automatic or inevitable (Papers II–III), and shaped by capacity rather than uniform across contexts (Paper V). These patterns are better aligned with the alternative vocabulary of *translation* where spread is not the frictionless propagation of a fixed object but the work of many actors who continuously interact with it [186, 187]. I retain “diffusion” to engage with the established tradition [38], but prefer “growth,” “deployment,” “adoption,” “uptake,” or “expansion” to better reflect the non-passivity of these processes. The point connects back to the notion of “unstoppable” or “inevitable” transitions which strip these processes of human agency, and risk trivialising the magnitude of socio-political effort that would be required to bring these transitions to fruition.

The language used to describe the transition is in turn closely linked to the tools used to analyse it. These typically include quantitative models of different kinds, and this thesis critically engages with techno-economic optimisation models. The mismatch between the technology trajectories produced by such models and empirically observed outcomes [14, 126, 149, 156, 188] can be attributed to the fact that such models are closed systems that do not contain all the generative mechanisms shaping growth in the real world. But this limitation is often lost in communication, especially to non-specialist audiences, and it is not difficult to mistake the models’ internal logic for reality, what Erica Thompson [189] calls staying inside *Model Land*.

In critiquing existing approaches, this thesis builds a model of its own (Paper III). In important respects PROLONG is designed to escape Model Land: it reports probabilities rather than point forecasts, it declines to estimate the deployment ceiling it cannot reliably know, it is validated against out-of-sample reality through hindcasting, and it is scoped to a purpose – medium-term growth rates – with its unsuitability for projecting saturation levels and for discontinuities stated

plainly. Yet it does not escape entirely, and it would be misleading (and ironic) to claim otherwise. PROLONG is trained on *simulated* ensembles built from assumed logistic and bilogistic forms; that training universe is itself a model world, and no amount of hindcasting can preempt real diffusion departing from the patterns the model is trained on. This places a structural limit on what the model can and cannot foresee. Its projection intervals, likewise, quantify the uncertainty the model can see, but not the deeper uncertainty of futures that are radically unlike the past and/or its training ensemble. This motivates the interpretation of PROLONG's projection as a probabilistic baseline for where existing trends might lead, conditional on not only past data but also the simulated worlds that inform the model.

5.4 Limitations

Section 5.3 addressed the limits of the kind of knowledge this thesis produces. The limitations set out here are of a narrower sort: they concern the scope of the cases studied, the data on which the analysis rests, the inferences the design can support, and the tools developed. The framework is built on, and applied to, technologies that succeeded in making it through the formative phase and spread widely. It is too early to empirically measure if technologies central to deep decarbonisation (e.g. electrolytic hydrogen, green steel), which are still in the formative phase also follow the same patterns. Much of the forward-looking analysis is also focused on the power sector. The framework and tools extend in principle to other technologies, as their application to mobile phones and internet use shows, but their suitability for use with technologies in sectors such as transport, industry, and heat remains to be demonstrated.

The data are uneven. This thesis relies on datasets assembled from deployment and cost records spanning two centuries and more than 200 countries, and their quality and coverage inevitably vary. They are typically sparser and less reliable for early historical series and for smaller and developing countries. The unavailability of certain kinds of data also constrains the types of questions this thesis can ask and confidently answer: the absence of R&D spending data restricts Paper IV to one-factor learning models. In Paper V, the overall penalty that later adopters pay is robust to controlling for national income. Disentangling the distinct components of absorptive capacity is harder,

because proxies such as national income and prior infrastructure are collinear and coverage is patchy. Absorptive capacity is therefore inferred from the combined effect of these proxies rather than isolated as a single measurable cause.

The regularities are also conditional on the contexts and periods that produced them. For instance, several of the conditions that shaped renewables growth over the last two decades are now shifting (e.g. higher interest rates, trade restrictions, the move from feed-in tariffs to auctions). Whether the phases and cruising speeds identified here persist in the future is an open question, and one that hindcasting over the coming years can begin to answer. Abrupt change is even harder to anticipate. Because the approach extrapolates from historical experience, it cannot anticipate discontinuities arising from geopolitical upheaval, radical technological breaks, or abrupt policy shifts. As argued in section 5.3, this is a limit of ontology rather than of data, and it is particularly important for PROLONG, whose training universe of simulated trajectories constrains the futures it can foresee.

Some limitations are specific to the tools developed. FOBI's continuous-growth criterion can misread technologies that deploy in large, discrete steps, such as nuclear power (Paper III); and DAIP's identification of inflection points and pulses can be sensitive to the choice of smoothing and persistence criteria (Paper II). Both FOBI and DAIP are designed for identifying takeoff, inflection points and pulses as they occur, but cannot anticipate them in advance. PROLONG projects deployment rates during the accelerating and steady growth phases rather than saturation levels, which limits what it can say about the developments in the formative phase, and about slowdown and saturation in the final phase (Paper III). It selects mature countries by the maturity of a fitted logistic curve, which Paper II shows disagrees with DAIP in a quarter to a third of countries, and its random forests cannot extrapolate beyond the logistic and bilogistic dynamics they are trained on.

5.5 Future research

These limitations, together with the avenues opened by the papers, point to three broad directions for future work: explaining the patterns identified here, extending data-driven projection to other technologies, sectors, and spatial scales, and widening the unit of analysis from

a single technology to the system it is embedded in.

The first direction is causal. Now that pulses can be detected, the natural next step is to explain why they occur when they do by disentangling the role of policy from market and societal forces across different contexts, and weighing the reinforcing policy–technology feedbacks that drive deployment upward against the countervailing ones that slow it down. The diagnostic toolkit developed here can support extensive inference, testing hypotheses about the effects of different policy instruments and contextual factors across many cases. Using an intensive rather than an extensive research design [181], the analysis of contrasting pairs, such as ones in which the same technology experiences pulses in one country but not in a comparable one, would let process-tracing follow the chain of cause and effect for particular pulses and recover the concrete causal history that an extensive design trades away. Such comparisons would also show whether a mechanism was necessary to induce a pulse, and would test the hypothesis that pulses are reactive sequences of opposition and response. Candidate mechanisms can also be built into simulation models to see whether they suffice to reproduce the steady growth phase and its pulses.

A related task is to make *policy-drivenness* itself measurable. Paper I treats it as a binary attribute of a technology, but also acknowledges that it more plausibly shifts along a continuum over the technology lifecycle. Quantifying the degree to which growth is policy-driven, and tracking how it changes as technologies mature, would sharpen the analysis of past growth, assessment of the overall policy effort expended on transitions, and the realism of projections. Extending the cost analysis in the same spirit, by incorporating R&D and other inputs where data permit, would allow the phase-structured decomposition of Paper IV to distinguish learning-by-doing from learning-by-researching [79, 190].

The second direction is to extend data-driven projection to other technologies, sectors and spatial scales so that it can play a more prominent role in modelling practice. That role requires projections beyond the power sector, for transport, industry, and heat, since soft-linking to IAMs and ESOMs needs coverage of the full portfolio of technologies those models represent. That role also requires three extensions to PROLONG itself, each following directly from its limitations. It could be made fully model-agnostic by removing the residual

assumptions about functional form in its training ensemble. It could be adapted to technologies that have taken off but not yet settled into steady growth. And its representation of the deployment ceiling could be improved. Because PROLONG projects growth rates rather than saturation levels, it can say little about how much a technology will ultimately grow. Two things are needed here: better diagnostics to distinguish a real approach towards saturation from the temporary decelerations that punctuate steady growth, and a way to estimate the ceiling itself *a priori*. Following Paper V, operationalising a country's absorptive capacity, resource potentials, and demand growth as variables conditioning national ceilings could enable PROLONG to project not only how *fast* but how *far* technologies will spread.

None of these extensions applies to technologies in the formative phase, where there are no consistent deployment trends to extrapolate from. Probabilistic assessment there can instead draw on recent approaches that use granular, project-level evidence on construction pipelines, announcements, and failure rates [137, 139, 159, 170, 171]. These belong to a broader family of probabilistic, data-driven projection methods that has emerged in recent years [27, 28, 191], each built around different data and questions. It is not yet clear how they compare in projection performance, or which questions each is best suited to answer. A systematic inter-comparison would help answer this. This would benchmark these methods against a common set of cases, with shared hindcasting protocols of the kind Paper III applies to PROLONG. The same benchmark could also admit mechanism-based simulation models alongside the data-driven methods. Run on shared cases, it would test empirically what Chapter 2 could only frame as a choice between an inside view, which builds mechanisms into models, and an outside view, which reasons from what comparable cases did (section 2.4) [12]. The comparison would give practitioners a basis for choosing among them, and lay the foundation for their coordinated use in integrated assessment and energy system modelling.

The third research direction would widen the unit of analysis from a single technology to the system in which it is embedded, in two respects. The first is the interdependence of new technologies. Deep decarbonisation requires the near-simultaneous diffusion of many interacting technologies including generation, storage, grids, heat pumps, electrified transport and industry – which compete for the same finance, infrastructure, and political attention in some settings

and reinforce one another in others [98, 192–196]. An empirically grounded understanding of these interdependencies would allow projections for a portfolio of technologies in a given system rather than one technology at a time. The second is the coupling between growth and decline. At the system level, the growth of low-carbon technologies is inseparable from the phase-out of the incumbents they displace – coal plants, internal-combustion vehicles, gas boilers, fossil-based industry – and the two processes are often coupled. They unfold within the same systems and compete for the same resources [197–199]. The diagnostic tools developed here could be adapted to analyse decline, asking whether phase-out follows its own characteristic phases and pulses, and how policy effort shifts between accelerating the new and retiring the old. Understanding that coupling – the sequencing and coordination of build-up and phase-down – is central to designing transitions that are not only fast but politically feasible.

CHAPTER 6

Conclusion

When Scientific American reported on Charles Brush's wind dynamo in 1890, it warned its readers that the cost of the machine more than offset the cheapness of the wind [1]. That verdict held for more than a century: wind and solar power cost too much to be considered a realistic alternative to other sources. It no longer does. The technologies are now cheap, and the question has shifted from whether the world can decarbonise its energy to how fast it will. Chapter 1 called the fall in cost only the first act in the energy story of our times. This thesis has been about the acts that follow: how technologies grow, how that growth can be measured while it is under way, and how far ahead its speed can be anticipated.

Growth follows recurring patterns, but it is far from predetermined (RQ1). Technologies do not trace a smooth S-curve to a ceiling. They move through four phases, each shaped by a balance of drivers and barriers that shifts as deployment and policy act on one another. For wind and solar, initial acceleration ends at about 3% of national electricity generation, far earlier than a logistic curve would place it. What follows is not decline but prolonged steady growth, punctuated by pulses of slowing and renewed growth (RQ1a). Consumer technologies such as mobile phones show the same pattern but reach their first inflection far later, a difference that plausibly reflects the larger role of policy in the growth of renewables. Costs follow the phases too: learning per doubling of deployment roughly doubles as technologies mature, yet annual cost reductions slow, because each doubling takes longer once growth stops compounding (RQ1b). And later adopters move through the lifecycle faster but grow more slowly and saturate lower, with ceilings that appear to depend more on a country's capacity to absorb a technology than on the timing of adoption (RQ1c).

This thesis introduces tools which make growth measurable while it is under way (RQ2). FOBI and DAIP identify takeoff, the first inflection point and the pulses that follow, directly from deployment data, without assuming a curve or knowing its ceiling. They are diagnostic rather than predictive: they show which phase a technology has entered, but cannot anticipate when it will move on to the next one. The peak growth rate at the first inflection point is a typical upper bound on the cruising speed that follows, while estimated ceilings and year-on-year growth rates, the metrics most often extrapolated, are the least reliable.

These regularities also make it possible to probabilistically project growth, for as long as the future resembles the past in structure (RQ3). PROLONG turns the experience of countries already further along in their trajectories into probabilistic projections for the world. Across hindcasts made from different points in the past, PROLONG's median projections came consistently closer to what happened than contemporaneous baseline scenarios of the IEA and the IPCC, which ran too low. Looking ahead, the central range from PROLONG puts global wind and solar growth in line with 2 °C-compatible pathways. A gap emerges against more ambitious targets. The COP28 pledge to triple renewable capacity by 2030 sits near PROLONG's 95th percentile, and meeting it would require wind growth in major economies to accelerate by 1.4 to 3 times and solar growth by 2 to 5 times. The 1.5 °C pathways with little or no overshoot exceed that percentile in the early 2030s, especially for solar. These conditional projections of where current dynamics lead act as a baseline for assessing ambition: a way of asking whether a target reflects what history suggests is achievable, or demands effort the world has rarely sustained.

Together these findings replace one picture of growth with another. In the old picture, a cheap technology diffuses on its own: passively, like a substance spreading down a gradient; automatically, like a contagion; and outward, from centre to periphery. History suggests otherwise. Growth is driven by a shifting combination of policy, markets and public response, not by cost alone. It pulses rather than proceeds automatically, and the effort behind it does not fall as a technology matures, because the barriers shift from cost and viability to integration, opposition and finance. And it does not flow evenly outward: latecomers inherit technology, but not the capacity to deploy it at scale. The decline in cost has removed an important obstacle to the trans-

ition, but many remain. The acceleration that the most ambitious pathways call for will depend on sustained policy where deployment is already under way, and on investment in institutions and absorptive capacity where they are weakest.

The speed at which the world deploys clean technology is not, in the end, a story about megawatts or S-curves. Technology is only ever a means. The end, as I see it, is the expansion of human capability, the freedom of people to live lives they have reason to value [200], now and for the generations who will inherit the climate we leave them. The skills, knowledge and institutions that let a country absorb new technology are themselves among those capabilities. Growth that is fast but uneven, that reaches some and passes others by, does not serve that end.

These questions reach beyond the energy transition, and beyond this thesis. I began it during a pandemic that changed daily life almost overnight, and finish it amid a boom in artificial intelligence, a technology described as a force that will remake the world whether or not anyone chooses it. I set out to ask how far ahead the growth of technology can be anticipated, and the honest answer is: only so far. That uncertainty is usually treated as a problem to be modelled away. I have come to think of it as the wiggle room in which societies' choices still matter: in the design of policy, the building of institutions, and the question of who benefits.

References

- [1] Mr. Brush's Windmill Dynamo (Dec. 1890). *Scientific American* **LXIII** (25), p. 389.
- [2] Ember (2026). *Global electricity review 2026*. Tech. rep. URL: <https://ember-energy.org/latest-insights/global-electricity-review-2026/> (Retrieved: 2026-09-15).
- [3] IEA (2025). *World energy outlook 2025*.
- [4] F. Creutzig, J. Hilaire, G. Nemet, F. Müller-Hansen and J. C. Minx (2023). Technological innovation enables low cost climate change mitigation. *Energy Research & Social Science* **105**, p. 103276. DOI: 10.1016/j.erss.2023.103276.
- [5] F. J. M. M. Nijse, J.-F. Mercure, N. Ameli, F. Larosa, S. Kothari, J. Rickman, P. Vercoulen and H. Pollitt (2023). The momentum of the solar energy transition. *Nature Communications* **14** (1), p. 6542. DOI: 10.1038/s41467-023-41971-7.
- [6] R. Way, M. C. Ives, P. Mealy and J. D. Farmer (2022). Empirically grounded technology forecasts and the energy transition. *Joule* **6** (9), pp. 2057–2082. DOI: 10.1016/j.joule.2022.08.009.
- [7] IRENA (2025). *Delivering on the UAE Consensus: tracking progress toward tripling renewable energy capacity and doubling energy efficiency by 2030*. Tech. rep. URL: <https://www.irena.org/Digital-Report/Tracking-progress-toward-tripling-renewable-energy-capacity-and-doubling-energy-efficiency-by-2030> (Retrieved: 2026-09-15).
- [8] K. Riahi, R. Schaeffer, J. Arango, K. Calvin, C. Guivarch, T. Hasegawa, K. Jiang, E. Kriegler, R. Matthews, G. Peters, A. Rao, S. Robertson, A. M. Sebbit, J. Steinberger, M. Tavoni and D. van Vuuren (2022). 'Mitigation pathways compatible with long-term goals – chapter 3'. In: *Climate change 2022: mitigation of climate change. Working Group III contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by P. R. Shukla, J. Skea, R. Slade, A. Al-Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas,

- R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz and J. Malley. DOI: 10.1017/9781009157926.005.
- [9] V. Krey, F. Guo, P. Kolp, W. Zhou, R. Schaeffer, A. Awasthy, C. Bertram, H.-S. de Boer, P. Fragkos, S. Fujimori, C. He, G. Iyer, K. Keramidas, A. C. Köberle, K. Oshiro, L. A. Reis, B. Shoai-Tehrani, S. Vishwanathan, P. Capros, L. Drouet, J. E. Edmonds, A. Garg, D. E. H. J. Gernaat, K. Jiang, M. Kannavou, A. Kitous, E. Kriegler, G. Luderer, R. Mathur, M. Muratori, F. Sano and D. P. van Vuuren (2019). Looking under the hood: a comparison of techno-economic assumptions across national and global integrated assessment models. *Energy* **172**, pp. 1254–1267. DOI: 10.1016/j.energy.2018.12.131.
- [10] F. W. Geels, F. Berkhout and D. P. van Vuuren (May 2016). Bridging analytical approaches for low-carbon transitions. *Nature Climate Change* **6** (6), pp. 576–583. DOI: 10.1038/nclimate2980.
- [11] E. Trutnevyte, L. F. Hirt, N. Bauer, A. Cherp, A. Hawkes, O. Y. Edelenbosch, S. Pedde and D. P. van Vuuren (2019). Societal transformations in models for energy and climate policy: the ambitious next step. *One Earth* **1** (4), pp. 423–433. DOI: 10.1016/j.oneear.2019.12.002.
- [12] J. Jewell and A. Cherp (2023). The feasibility of climate action: bridging the inside and the outside view through feasibility spaces. *WIREs Climate Change* **14** (5). DOI: 10.1002/wcc.838.
- [13] F. Creutzig, P. Agoston, J. C. Goldschmidt, G. Luderer, G. Nemet and R. C. Pietzcker (Aug. 2017). The underestimated potential of solar energy to mitigate climate change. *Nature Energy* **2** (9), p. 17140. DOI: 10.1038/nenergy.2017.140.
- [14] G. Lopez, Y. Pourjamal and C. Breyer (2025). Paving the way towards a sustainable future or lagging behind? an ex-post analysis of the international energy agency's world energy outlook. *Renewable and Sustainable Energy Reviews* **212**, p. 115371. DOI: 10.1016/j.rser.2025.115371.
- [15] A. Cherp, V. Vinichenko, J. Tosun, J. A. Gordon and J. Jewell (2021). National growth dynamics of wind and solar power compared to the growth required for global climate targets. *Nature Energy* **6** (7), pp. 742–754. DOI: 10.1038/s41560-021-00863-0.
- [16] V. Vinichenko, J. Jewell, J. Jacobsson and A. Cherp (2023). Historical diffusion of nuclear, wind and solar power in different national contexts: implications for climate mitigation pathways. *Environmental Research Letters* **18** (9), p. 094066. DOI: 10.1088/1748-9326/acf47a.

-
- [17] Z. Griliches (1957). Hybrid corn: an exploration in the economics of technological change. *Econometrica* **25** (4), p. 501. DOI: 10.2307/1905380.
- [18] C. Marchetti and N. Nakicenovic (1979). *The dynamics of energy systems and the logistic substitution model*. Tech. rep. Laxenburg; Austria: IIASA. URL: <https://pure.iiasa.ac.at/id/eprint/1024/>.
- [19] A. Grubler (1990). *The rise and fall of infrastructures dynamics of evolution and technological change*. Heidelberg; Germany: Physica-Verlag Heidelberg. URL: <https://pure.iiasa.ac.at/id/eprint/3351/>.
- [20] C. Wilson, A. Grubler, N. Bento, S. Healey, S. De Stercke and C. Zimm (2020). Granular technologies to accelerate decarbonization. *Science* **368** (6486), pp. 36–39. DOI: 10.1126/science.aaz8060.
- [21] J. P. Martino (2003). A review of selected recent advances in technological forecasting. *Technological Forecasting and Social Change* **70** (8), pp. 719–733. DOI: 10.1016/s0040-1625(02)00375-x.
- [22] J. P. Hansen, P. A. Narbel and D. L. Aksnes (Apr. 2017). Limits to growth in the renewable energy sector. *Renewable and Sustainable Energy Reviews* **70**, pp. 769–774. DOI: 10.1016/j.rser.2016.11.257.
- [23] D. N. Madsen and J. P. Hansen (2019). Outlook of solar energy in Europe based on economic growth characteristics. *Renewable and Sustainable Energy Reviews* **114**, p. 109306. DOI: 10.1016/j.rser.2019.109306.
- [24] C. L. Baumgärtner, R. Way, M. C. Ives and J. D. Farmer (2024). The need for better statistical testing in data-driven energy technology modeling. *Joule* **8** (9), pp. 2453–2466. DOI: 10.1016/j.joule.2024.07.016.
- [25] Z. Griliches (1980). Hybrid Corn revisited: a reply. *Econometrica* **48** (6), pp. 1463–1465. URL: <https://www.jstor.org/stable/1912818>.
- [26] A. Debecker and T. Modis (1994). Determination of the uncertainties in S-curve logistic fits. *Technological Forecasting and Social Change* **46** (2), pp. 153–173. DOI: 10.1016/0040-1625(94)90023-x.
- [27] N. Zielonka, X. Wen and E. Trutnevyte (2023). Probabilistic projections of granular energy technology diffusion at subnational level. *PNAS Nexus* **2** (10), pgad321. DOI: 10.1093/pnasnexus/pgad321.
- [28] N. Zielonka and E. Trutnevyte (2025). Probabilities of reaching required diffusion of granular energy technologies in European countries. *iScience* **28** (2), p. 111825. DOI: 10.1016/j.isci.2025.111825.

- [29] R. Fouquet (2010). The slow search for solutions: lessons from historical energy transitions by sector and service. *Energy Policy* **38** (11), pp. 6586–6596. DOI: 10.1016/j.enpol.2010.06.029.
- [30] V. Smil (2017). *Energy and civilization: a history*. The MIT Press. ISBN: 9780262035774. DOI: 10.7551/mitpress/9780262035774.001.0001.
- [31] A. Grubler (1996). Time for a change: on the patterns of diffusion of innovation. *Daedalus* **125** (3), pp. 19–42. DOI: 10.2307/20027369.
- [32] R. Fouquet and P. J. G. Pearson (Jan. 2012). Past and prospective energy transitions: insights from history. *Energy Policy* **50**, pp. 1–7. DOI: 10.1016/j.enpol.2012.08.014.
- [33] A. Grubler (1998). Technology and global change. *International Institute for Applied Systems Analysis, Austria*. DOI: 10.1017/cbo9781316036471.
- [34] A. Grubler, C. Wilson and G. Nemet (Jan. 2016). Apples, oranges, and consistent comparisons of the temporal dynamics of energy transitions. *Energy Research & Social Science* **22**, pp. 18–25. DOI: 10.1016/j.erss.2016.08.015.
- [35] J. Rogelj, T. Fransen, M. G. J. den Elzen, R. D. Lamboll, C. Schumer, T. Kuramochi, F. Hans, S. Mooldijk and J. Portugal-Pereira (2023). Credibility gap in net-zero climate targets leaves world at high risk. *Science* **380** (6649), pp. 1014–1016. DOI: 10.1126/science.adg6248.
- [36] T. W. Valente and E. M. Rogers (1995). The origins and development of the diffusion of innovations paradigm as an example of scientific growth. *Science Communication* **16** (3), pp. 242–273. DOI: 10.1177/1075547095016003002.
- [37] B. Ryan and N. C. Gross (1943). The diffusion of hybrid seed corn in two Iowa communities. *Rural Sociology* **8** (1), pp. 15–24. URL: <https://www.proquest.com/openview/7de2b2276a089fe888071663de12b6a0/1?cb1=1817355&pq-origsite=gscholar>.
- [38] E. M. Rogers (1983). *Diffusion of innovations*. 3rd ed. New York; United States: The Free Press.
- [39] R. M. Solow (1957). Technical change and the aggregate production function. *The Review of Economics and Statistics* **39** (3), pp. 312–320.
- [40] K. J. Arrow (1962). The economic implications of learning by doing. *The Review of Economic Studies* **29** (3), pp. 155–173. DOI: 10.2307/2295952.

-
- [41] R. R. Nelson and S. G. Winter (1982). *An evolutionary theory of economic change*. Harvard University Press.
- [42] G. Dosi (1982). Technological paradigms and technological trajectories. *Research Policy* **11** (3), pp. 147–162.
- [43] L. Brandner and M. A. Straus (1959). Congruence versus profitability in the diffusion of hybrid sorghum. *Rural Sociology* **24** (4), pp. 381–383.
- [44] E. M. Rogers and A. Havens (1962). Predicting innovativeness*. *Sociological Inquiry* **32** (1), pp. 34–42. DOI: 10.1111/j.1475-682x.1962.tb00528.x.
- [45] J. M. Love (1980). ‘Does a new technology’s profitability affect its diffusion?’ Cornell University, Charles H. Dyson School of Applied Economics and Management, Ithaca, NY.
- [46] E. Mansfield (1961). Technical change and the rate of imitation. *Econometrica* **29** (4), pp. 741–766. URL: <https://www.jstor.org/stable/1911817?origin=crossref>.
- [47] J. C. Fisher and R. H. Pry (1971). A simple substitution model of technological change. *Technological Forecasting and Social Change* **3**, pp. 75–88. DOI: 10.1016/s0040-1625(71)80005-7.
- [48] A. Grubler, N. Nakicenovic and D. G. Victor (May 1999). Dynamics of energy technologies and global change. *Energy Policy* **27** (5), pp. 247–280. DOI: 10.1016/s0301-4215(98)00067-6.
- [49] C. Wilson and A. Grubler (2011). Lessons from the history of technological change for clean energy scenarios and policies. *Natural Resources Forum* **35** (3), pp. 165–184. DOI: 10.1111/j.1477-8947.2011.01386.x.
- [50] C. Wilson, A. Grubler, N. Bauer, V. Krey and K. Riahi (May 2013). Future capacity growth of energy technologies: are scenarios consistent with historical evidence? *Climatic Change* **118** (2), pp. 381–395. DOI: 10.1007/s10584-012-0618-y.
- [51] F. M. Bass (1969). A new product growth for model consumer durables. *Management Science* **15** (5), pp. 215–227.
- [52] R. Peres, E. Muller and V. Mahajan (2010). Innovation diffusion and new product growth models: a critical review and research directions. *International Journal of Research in Marketing* **27** (2), pp. 91–106. DOI: 10.1016/j.ijresmar.2009.12.012.
- [53] A. Gerschenkron (1962). *Economic backwardness in historical perspective*. Harvard University Press.

- [54] J. Goldemberg (Jan. 1998). Leapfrog energy technologies. *Energy Policy* **26** (10), pp. 729–741. DOI: 10 . 1016 / s0301 - 4215 (98) 00025-1.
- [55] D. Comin and B. Hobijn (2004). Cross-country technology adoption: making the theories face the facts. *Journal of Monetary Economics* **51** (1), pp. 39–83. DOI: 10 . 1016 / j . jmoneco . 2003 . 07 . 003.
- [56] D. Comin and M. Mestieri (2018). If technology has arrived everywhere, why has income diverged? *American Economic Journal: Macroeconomics* **10** (3), pp. 137–178. DOI: 10 . 1257 / mac . 20150175.
- [57] J. H. Greene, M. J. Gidden, E. Brutschin and G. F. Nemet (2026). Drivers of technology diffusion speed in countries. *Nature Communications* **17** (1), p. 7591. DOI: 10 . 1038 / s41467 - 026 - 73563 - 6.
- [58] N. Bento and C. Wilson (Jan. 2016). Measuring the duration of formative phases for energy technologies. *Environmental Innovation and Societal Transitions* **21**, pp. 95–112. DOI: 10 . 1016 / j . eist . 2016 . 04 . 004.
- [59] N. Bento, C. Wilson and L. D. Anadon (Jan. 2018). Time to get ready: conceptualizing the temporal and spatial dynamics of formative phases for energy technologies. *Energy Policy* **119**, pp. 282–293. DOI: 10 . 1016 / j . enpol . 2018 . 04 . 015.
- [60] J. Rotmans, R. Kemp and M. van Asselt (2001). More evolution than revolution: transition management in public policy. *foresight* **3** (1), pp. 15–31. DOI: 10 . 1108 / 14636680110803003.
- [61] J. Markard (Jan. 2018). The next phase of the energy transition and its implications for research and policy. *Nature Energy* **3** (8), pp. 628–633. DOI: 10 . 1038 / s41560 - 018 - 0171 - 7.
- [62] J. Gosens, F. Hedenus and B. A. Sandén (2017). Faster market growth of wind and PV in late adopters due to global experience build-up. *Energy* **131**, pp. 267–278. DOI: 10 . 1016 / j . energy . 2017 . 05 . 046.
- [63] P. F. Verhulst (1845). Recherches mathématiques sur la loi d'accroissement de la population. *Nouveaux Mémoires de l'Académie Royale des Sciences et Belles-Lettres de Bruxelles* **18** (1), pp. 1–45.
- [64] R. Pearl and L. J. Reed (1920). On the rate of growth of the population of the United States since 1790 and its mathematical representation. *Proceedings of the National Academy of Sciences* **6** (6), pp. 275–288. DOI: 10 . 1073 / pnas . 6 . 6 . 275.
- [65] C. Marchetti (1983). The automobile in a system context. *Technological Forecasting and Social Change* **23** (1), pp. 3–23. DOI: 10 . 1016 / 0040 - 1625 (83) 90068 - 9.

-
- [66] M. Harris (1968). *The rise of anthropological theory: a history of theories of culture*. Thomas Y. Crowell Company.
- [67] T. P. Wright (1936). Factors affecting the cost of airplanes. *Journal of the aeronautical sciences* **3** (4), pp. 122–128.
- [68] G. E. Moore (1998). Cramming more components onto integrated circuits. *Proceedings of the IEEE* **86**, pp. 82–85.
- [69] B. Nagy, J. D. Farmer, Q. M. Bui and J. E. Trancik (Jan. 2013). Statistical basis for predicting technological progress. *PLoS ONE* **8** (2), e52669. DOI: 10.1371/journal.pone.0052669.
- [70] W. D. Nordhaus (2014). The perils of the learning model for modeling endogenous technological change. *The Energy Journal* **35** (1), pp. 1–13. URL: <https://www.jstor.org/stable/24693815>.
- [71] A. McDonald and L. Schrattenholzer (2002). Learning curves and technology assessment. *International Journal of Technology Management* **23** (7–8), pp. 718–745. DOI: 10.1504/ijtm.2002.003035.
- [72] E. S. Rubin, I. M. L. Azevedo, P. Jaramillo and S. Yeh (2015). A review of learning rates for electricity supply technologies. *Energy Policy* **86**, pp. 198–218. DOI: 10.1016/j.enpol.2015.06.011.
- [73] G. F. Nemet (2006). Beyond the learning curve: factors influencing cost reductions in photovoltaics. *Energy Policy* **34** (17), pp. 3218–3232. DOI: 10.1016/j.enpol.2005.06.020.
- [74] F. Ferioli, K. Schoots and B. C. C. van der Zwaan (2009). Use and limitations of learning curves for energy technology policy: a component-learning hypothesis. *Energy Policy* **37** (7), pp. 2525–2535. DOI: 10.1016/j.enpol.2008.10.043.
- [75] G. Kavlak, J. McNerney and J. E. Trancik (2018). Evaluating the causes of cost reduction in photovoltaic modules. *Energy Policy* **123**, pp. 700–710. DOI: 10.1016/j.enpol.2018.08.015.
- [76] A. Malhotra and T. S. Schmidt (2020). Accelerating low-carbon innovation. *Joule* **4** (11), pp. 2259–2267. DOI: 10.1016/j.joule.2020.09.004.
- [77] J. D. Farmer and F. Lafond (2016). How predictable is technological progress? *Research Policy* **45** (3), pp. 647–665. DOI: 10.1016/j.respol.2015.11.001.
- [78] J. Behrens, E. Zeyen, M. Hoffmann, D. Stolten and J. M. Weinand (2024). Reviewing the complexity of endogenous technological learning for energy system modeling. *Advances in Applied Energy* **16**, p. 100192. DOI: 10.1016/j.adapen.2024.100192.

- [79] S. Yeh and E. S. Rubin (2012). A review of uncertainties in technology experience curves. *Energy Economics* **34** (3), pp. 762–771. DOI: 10 . 1016/j . eneco . 2011 . 11 . 006.
- [80] M. Wei, S. J. Smith and M. D. Sohn (2017). Non-constant learning rates in retrospective experience curve analyses and their correlation to deployment programs. *Energy Policy* **107**, pp. 356–369. DOI: 10 . 1016/j . enpol . 2017 . 04 . 035.
- [81] A. Carlino, A. Wongel, L. Duan, E. Virgüez, S. J. Davis, M. R. Edwards and K. Caldeira (2025). Variability of technology learning rates. *Advances in Applied Energy* **20**, p. 100252. DOI: 10 . 1016/j . adapen . 2025 . 100252.
- [82] R. Kemp, J. Schot and R. Hoogma (1998). Regime shifts to sustainability through processes of niche formation: the approach of strategic niche management. *Technology Analysis & Strategic Management* **10** (2), pp. 175–198. DOI: 10 . 1080/09537329808524310.
- [83] F. W. Geels (2002). Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. *Research Policy* **31** (8-9), pp. 1257–1274. DOI: 10 . 1016 / s0048 - 7333(02)00062-8.
- [84] S. Jacobsson and A. Bergek (Oct. 2004). Transforming the energy sector: the evolution of technological systems in renewable energy technology. *Industrial and Corporate Change* **13** (5), pp. 815–849. DOI: 10 . 1093/icc/dth032.
- [85] M. P. Hekkert, R. A. A. Suurs, S. O. Negro, S. Kuhlmann and R. E. H. M. Smits (2007). Functions of innovation systems: a new approach for analysing technological change. *Technological Forecasting and Social Change* **74** (4), pp. 413–432. DOI: 10 . 1016/j . techfore . 2006 . 03 . 002.
- [86] A. Bergek, S. Jacobsson, B. Carlsson, S. Lindmark and A. Rickne (Jan. 2008). Analyzing the functional dynamics of technological innovation systems: a scheme of analysis. *Research Policy* **37** (3), pp. 407–429. DOI: 10 . 1016/j . respol . 2007 . 12 . 003.
- [87] F. W. Geels (2014). Regime resistance against low-carbon transitions: introducing politics and power into the multi-level perspective. *Theory, Culture & Society* **31** (5), pp. 21–40. DOI: 10 . 1177 / 0263276414531627.
- [88] F. Geels and R. Raven (2006). Non-linearity and expectations in niche-development trajectories: ups and downs in Dutch biogas development (1973–2003). *Technology Analysis & Strategic Management* **18** (3-4), pp. 375–392. DOI: 10 . 1080/09537320600777143.

-
- [89] V. Lauber and S. Jacobsson (Mar. 2016). The politics and economics of constructing, contesting and restricting socio-political space for renewables – the German Renewable Energy Act. *Environmental Innovation and Societal Transitions* **18**, pp. 147–163. DOI: 10.1016/j.eist.2015.06.005.
- [90] A. Cherp, V. Vinichenko, J. Jewell, E. Brutschin and B. K. Sovacool (2018). Integrating techno-economic, socio-technical and political perspectives on national energy transitions: a meta-theoretical framework. *Energy Research & Social Science* **37**, pp. 175–190. DOI: 10.1016/j.erss.2017.09.015.
- [91] F. Kern and K. S. Rogge (2018). Harnessing theories of the policy process for analysing the politics of sustainability transitions: a critical survey. *Environmental Innovation and Societal Transitions* **27**, pp. 102–117. DOI: 10.1016/j.eist.2017.11.001.
- [92] D. Béland and E. Schlager (2019). Varieties of policy feedback research: looking backward, moving forward. *Policy Studies Journal* **47** (2), pp. 184–205. DOI: 10.1111/psj.12340.
- [93] M. Lockwood (2022). Policy feedback and institutional context in energy transitions. *Policy Sciences* **55** (3), pp. 487–507. DOI: 10.1007/s11077-022-09467-1.
- [94] S. Schaub, J. Tosun and A. J. Jordan (2024). Climate action through policy expansion and/or dismantling: country-comparative insights: an introduction to the special issue. *Journal of Comparative Policy Analysis: Research and Practice* **26** (3-4), pp. 215–232. DOI: 10.1080/13876988.2024.2369640.
- [95] L. C. Stokes (2016). Electoral backlash against climate policy: a natural experiment on retrospective voting and local resistance to public policy. *American Journal of Political Science* **60** (4), pp. 958–974. DOI: 10.1111/ajps.12220.
- [96] J. J. Patterson (2023). Backlash to climate policy. *Global Environmental Politics* **23** (1), pp. 68–90. DOI: 10.1162/glep_a_00684.
- [97] D. L. Edmondson, F. Kern and K. S. Rogge (2019). The co-evolution of policy mixes and socio-technical systems: towards a conceptual framework of policy mix feedback in sustainability transitions. *Research Policy* **48** (10), p. 103555. DOI: 10.1016/j.respol.2018.03.010.
- [98] F. W. Geels and M. Ayoub (2025). Policy feedbacks and socio-technical feedbacks in accelerated low-carbon transitions: an integrated conceptual framework illustrated with case studies of lighting and smart meters. *Research Policy* **54** (8), p. 105294. DOI: 10.1016/j.respol.2025.105294.

- [99] W. B. Arthur (Mar. 1989). Competing technologies, increasing returns, and lock-in by historical small events. *The Economic Journal* **99** (394), pp. 116–131. DOI: 10.2307/2234208.
- [100] J. Meckling, N. Kelsey, E. Biber and J. Zysman (2015). Winning coalitions for climate policy. *Science* **349** (6253), pp. 1170–1171. DOI: 10.1126/science.aab1336.
- [101] T. S. Schmidt and S. Sewerin (2017). Technology as a driver of climate and energy politics. *Nature Energy* **2** (6), p. 17084. DOI: 10.1038/nenergy.2017.84.
- [102] M. Pahle, D. Burtraw, C. Flachsland, N. Kelsey, E. Biber, J. Meckling, O. Edenhofer and J. Zysman (Jan. 2018). Sequencing to ratchet up climate policy stringency. *Nature Climate Change* **8** (10), pp. 861–867. DOI: 10.1038/s41558-018-0287-6.
- [103] G. C. Unruh (Jan. 2000). Understanding carbon lock-in. *Energy Policy* **28** (12), pp. 817–830. DOI: 10.1016/s0301-4215(00)00070-7.
- [104] L. Susskind, J. Chun, A. Gant, C. Hodgkins, J. Cohen and S. Lohmar (2022). Sources of opposition to renewable energy projects in the United States. *Energy Policy* **165**, p. 112922. DOI: 10.1016/j.enpol.2022.112922.
- [105] L. C. Stokes, E. Franzblau, J. R. Lovering and C. Miljanich (2023). Prevalence and predictors of wind energy opposition in North America. *Proceedings of the National Academy of Sciences* **120** (40), e2302313120. DOI: 10.1073/pnas.2302313120.
- [106] P. J. Heptonstall and R. J. K. Gross (2021). A systematic review of the costs and impacts of integrating variable renewables into power grids. *Nature Energy* **6** (1), pp. 72–83. DOI: 10.1038/s41560-020-00695-4.
- [107] B. Frantál, M. Frolova and J. Liñán-Chacón (2023). Conceptualizing the patterns of land use conflicts in wind energy development: towards a typology and implications for practice. *Energy Research & Social Science* **95**, p. 102907. DOI: 10.1016/j.erss.2022.102907.
- [108] F. W. Geels (2022). Causality and explanation in socio-technical transitions research: mobilising epistemological insights from the wider social sciences. *Research Policy* **51** (6), p. 104537. DOI: 10.1016/j.respol.2022.104537.
- [109] H. Breetz, M. Mildenberger and L. Stokes (2018). The political logics of clean energy transitions. *Business and Politics* **20** (4), pp. 492–522. DOI: 10.1017/bap.2018.14.
- [110] I. D. Smith (2020). How the process of transitions shapes the politics of decarbonization: tracing policy feedback effects across phases of

-
- the energy transition. *Energy Research & Social Science* **70**, p. 101753. DOI: 10.1016/j.erss.2020.101753.
- [111] L. F. Hirt, G. Schell, M. Sahakian and E. Trutnevyte (2020). A review of linking models and socio-technical transitions theories for energy and climate solutions. *Environmental Innovation and Societal Transitions* **35**, pp. 162–179. DOI: 10.1016/j.eist.2020.03.002.
 - [112] S. Sorrell (2018). Explaining sociotechnical transitions: a critical realist perspective. *Research Policy* **47** (7), pp. 1267–1282. DOI: 10.1016/j.respol.2018.04.008.
 - [113] L. van Beek, M. Hajer, P. Pelzer, D. van Vuuren and C. Cassen (2020). Anticipating futures through models: the rise of integrated assessment modelling in the climate science-policy interface since 1970. *Global Environmental Change* **65**, p. 102191. DOI: 10.1016/j.gloenvcha.2020.102191.
 - [114] D. H. Meadows, D. L. Meadows, J. Randers and W. W. Behrens (1972). *The limits to growth: a report for the club of rome's project on the predicament of mankind*.
 - [115] W. D. Nordhaus (1973). The allocation of energy resources. *Brookings Papers on Economic Activity* **4** (3), pp. 529–576. URL: https://www.brookings.edu/wp-content/uploads/1973/12/1973c_bpe_a_nordhaus_houthakker_solow.pdf.
 - [116] IIASA (1981). *Energy in a finite world: report*. Tech. rep. Cambridge, Massachusetts: Ballinger Publishing Company.
 - [117] W. D. Nordhaus (1977). Economic growth and climate: the carbon dioxide problem. *The American Economic Review*. Papers and Proceedings of the Eighty-ninth Annual Meeting of the American Economic Association **67** (1), pp. 341–346. URL: <https://www.jstor.org/stable/1815926>.
 - [118] — (1993). Rolling the ‘DICE’: an optimal transition path for controlling greenhouse gases. *Resource and Energy Economics* **15** (1), pp. 27–50. DOI: 10.1016/0928-7655(93)90017-o.
 - [119] S. Pfenninger, A. Hawkes and J. Keirstead (2014). Energy systems modeling for twenty-first century energy challenges. *Renewable and Sustainable Energy Reviews* **33**, pp. 74–86. DOI: 10.1016/j.rser.2014.02.003.
 - [120] K. Anderson and J. Jewell (2019). Debating the bedrock of climate-change mitigation scenarios. *Nature* **573** (7774), pp. 348–349. DOI: 10.1038/d41586-019-02744-9.
 - [121] L. van Beek, J. Oomen, M. Hajer, P. Pelzer and D. van Vuuren (2022). Navigating the political: an analysis of political calibration of integ-

- rated assessment modelling in light of the 1.5 °C goal. *Environmental Science & Policy* **133**, pp. 193–202. DOI: 10.1016/j.envsci.2022.03.024.
- [122] B. Cointe and H. Guillemot (2023). A history of the 1.5°C target. *Wiley Interdisciplinary Reviews: Climate Change*. DOI: 10.1002/wcc.824.
- [123] I. Keppo, I. Butnar, N. Bauer, M. Caspani, O. Edelenbosch, J. Emmerling, P. Fragkos, C. Guivarch, M. Harmsen, J. Lefèvre, T. Le Gallic, M. Leimbach, W. McDowall, J.-F. Mercure, R. Schaeffer, E. Trutnevyte and F. Wagner (2021). Exploring the possibility space: taking stock of the diverse capabilities and gaps in integrated assessment models. *Environmental Research Letters* **16** (5), p. 053006. DOI: 10.1088/1748-9326/abe5d8.
- [124] S. Messner (1997). Endogenized technological learning in an energy systems model. *Journal of Evolutionary Economics* **7** (3), pp. 291–313. DOI: 10.1007/s001910050045.
- [125] L. Steg, J. Veldstra, K. de Kleijne, Ş. Kılış, A. F. P. Lucena, L. J. Nilsson, M. Sugiyama, P. Smith, M. Tavoni, H. de Coninck, R. van Diemen, P. Renforth, S. Mirasgedis, G. Nemet, R. Görsch, H. Muri, P. Bertoldi, L. F. Cabeza, É. Mata, A. Novikova, L. R. Caldas, M. Chàfer, R. Khosla and D. Vérez (2022). A method to identify barriers to and enablers of implementing climate change mitigation options. *One Earth* **5** (11), pp. 1216–1227. DOI: 10.1016/j.oneear.2022.10.007.
- [126] V. Fisch-Romito, M. Jaxa-Rozen, X. Wen and E. Trutnevyte (2025). Multi-country evidence on societal factors to include in energy transition modelling. *Nature Energy*, pp. 1–10. DOI: 10.1038/s41560-025-01719-7.
- [127] S. L. Bi, N. Bauer and J. Jewell (2023). Coal-exit alliance must confront freeriding sectors to propel Paris-aligned momentum. *Nature Climate Change* **13** (2), pp. 130–139. DOI: 10.1038/s41558-022-01570-8.
- [128] A. Krumm, D. Süsner and P. Blechinger (2022). Modelling social aspects of the energy transition: what is the current representation of social factors in energy models? *Energy* **239**, p. 121706. DOI: 10.1016/j.energy.2021.121706.
- [129] N. Nakicenovic, J. Alcamo, G. Davis, B. de Vries, J. Fenhann, S. Gaffin, K. Gregory, A. Grubler, T. Y. June, T. Kram, E. L. La Rovere, L. Michaelis, S. Mori, T. Morita, W. Pepper, H. Pitcher, L. Price, K. Riahi, A. Roehrl, H.-H. Rogner, A. Sankovski, M. Schlesinger, P. Shukla, S. Smith, R. Swart, S. van Rooijen, N. Victor and Z. Dadi (2000). *Special report on emissions scenarios*. Cambridge University Press.

-
- [130] B. C. O'Neill, E. Kriegler, K. Riahi, K. L. Ebi, S. Hallegatte, T. R. Carter, R. Mathur and D. P. van Vuuren (2014). A new scenario framework for climate change research: the concept of shared socioeconomic pathways. *Climatic Change* **122** (3), pp. 387–400. DOI: 10.1007/s10584-013-0905-2.
 - [131] K. Riahi, D. P. van Vuuren, E. Kriegler, J. Edmonds, B. C. O'Neill, S. Fujimori, N. Bauer, K. Calvin, R. Dellink, O. Fricko, W. Lutz, A. Popp, J. C. Cuaresma, S. KC, M. Leimbach, L. Jiang, T. Kram, S. Rao, J. Emmerling, K. Ebi, T. Hasegawa, P. Havlik, F. Humpenöder, L. Aleluia Da Silva, S. Smith, E. Stehfest, V. Bosetti, J. Eom, D. Gernaat, T. Masui, J. Rogelj, J. Strefler, L. Drouet, V. Krey, G. Luderer, M. Harmsen, K. Takahashi, L. Baumstark, J. C. Doelman, M. Kainuma, Z. Klimont, G. Marangoni, H. Lotze-Campen, M. Obersteiner, A. Tabeau and M. Tavoni (Jan. 2017). The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: an overview. *Global Environmental Change* **42**, pp. 153–168. DOI: 10.1016/j.gloenvcha.2016.05.009.
 - [132] B. D. Leibowicz, V. Krey and A. Grubler (2016). Representing spatial technology diffusion in an energy system optimization model. *Technological Forecasting and Social Change* **103**, pp. 350–363. DOI: 10.1016/j.techfore.2015.06.001.
 - [133] M. G. Morgan and D. W. Keith (2008). Improving the way we think about projecting future energy use and emissions of carbon dioxide. *Climatic Change* **90** (3), pp. 189–215. DOI: 10.1007/s10584-008-9458-1.
 - [134] M. Jaxa-Rozen and E. Trutnevyte (2021). Sources of uncertainty in long-term global scenarios of solar photovoltaic technology. *Nature Climate Change* **11** (3), pp. 266–273. DOI: 10.1038/s41558-021-00998-8.
 - [135] M. Victoria, N. Haegel, I. M. Peters, R. Sinton, A. Jäger-Waldau, C. del Cañizo, C. Breyer, M. Stocks, A. Blakers, I. Kaizuka, K. Komoto and A. Smets (2021). Solar photovoltaics is ready to power a sustainable future. *Joule* **5** (5), pp. 1041–1056. DOI: 10.1016/j.joule.2021.03.005.
 - [136] G. F. Nemet, M. J. Gidden, J. Greene, C. Roberts, W. F. Lamb, J. C. Minx, S. M. Smith, O. Geden and K. Riahi (2023). Near-term deployment of novel carbon removal to facilitate longer-term deployment. *Joule* **7** (12). DOI: 10.1016/j.joule.2023.11.001.
 - [137] T. Kazlou, A. Cherp and J. Jewell (2024). Feasible deployment of carbon capture and storage and the requirements of climate targets.

- Nature Climate Change* **14** (10), pp. 1047–1055. DOI: 10 . 1038 / s 41558-024-02104-0.
- [138] A. Odenweller, F. Ueckerdt, G. F. Nemet, M. Jensterle and G. Luderer (2022). Probabilistic feasibility space of scaling up green hydrogen supply. *Nature Energy* **7** (9), pp. 854–865. DOI: 10 . 1038 / s41560-022-01097-4.
- [139] A. Odenweller and F. Ueckerdt (2025). The green hydrogen ambition and implementation gap. *Nature Energy* **10** (1), pp. 110–123. DOI: 10 . 1038 / s41560-024-01684-7.
- [140] O. Edenhofer and M. Kowarsch (2015). Cartography of pathways: a new model for environmental policy assessments. *Environmental Science & Policy* **51**, pp. 56–64. DOI: 10 . 1016 / j . envsci . 2015 . 03 . 017.
- [141] M. Haasnoot, R. Biesbroek, J. Lawrence, V. Muccione, R. Lempert and B. Glavovic (2020). Defining the solution space to accelerate climate change adaptation. *Regional Environmental Change* **20** (2), p. 37. DOI: 10 . 1007 / s10113-020-01623-8.
- [142] K. Riahi, E. Kriegler, N. Johnson, C. Bertram, M. den Elzen, J. Eom, M. Schaeffer, J. Edmonds, M. Isaac, V. Krey, T. Longden, G. Luderer, A. Méjean, D. L. McCollum, S. Mima, H. Turton, D. P. van Vuuren, K. Wada, V. Bosetti, P. Capros, P. Criqui, M. Hamdi-Cherif, M. Kainuma and O. Edenhofer (2015). Locked into Copenhagen pledges — implications of short-term emission targets for the cost and feasibility of long-term climate goals. *Technological Forecasting and Social Change* **90**, pp. 8–23. DOI: 10 . 1016 / j . techfore . 2013 . 09 . 016.
- [143] E. Brutschin, S. Pianta, M. Tavoni, K. Riahi, V. Bosetti, G. Marangoni and B. J. van Ruijven (June 2021). A multidimensional feasibility evaluation of low-carbon scenarios. *Environmental Research Letters* **16** (6), p. 064069. DOI: 10 . 1088 / 1748-9326 / abf0ce.
- [144] F. W. Geels, A. McMeekin and B. Pfluger (2020). Socio-technical scenarios as a methodological tool to explore social and political feasibility in low-carbon transitions: bridging computer models and the multi-level perspective in UK electricity generation (2010–2050). *Technological Forecasting and Social Change* **151**, p. 119258. DOI: 10 . 1016 / j . techfore . 2018 . 04 . 001.
- [145] K. S. Rogge, B. Pfluger and F. W. Geels (2020). Transformative policy mixes in socio-technical scenarios: the case of the low-carbon transition of the German electricity system (2010–2050). *Technological Forecasting and Social Change* **151**, p. 119259. DOI: 10 . 1016 / j . techfore . 2018 . 04 . 002.

-
- [146] D. Kahneman and D. Lovallo (1993). Timid choices and bold forecasts: a cognitive perspective on risk taking. *Management Science* **39** (1), pp. 17–31. DOI: 10.1287/mnsc.39.1.17.
 - [147] G. Holtz, F. Alkemade, F. de Haan, J. Köhler, E. Trutnevyte, T. Luthe, J. Halbe, G. Papachristos, E. Chappin, J. Kwakkel and S. Ruutu (2015). Prospects of modelling societal transitions: position paper of an emerging community. *Environmental Innovation and Societal Transitions* **17**, pp. 41–58. DOI: 10.1016/j.eist.2015.05.006.
 - [148] J.-F. Mercure, H. Pollitt, A. M. Bassi, J. E. Viñuales and N. R. Edwards (2016). Modelling complex systems of heterogeneous agents to better design sustainability transitions policy. *Global Environmental Change* **37**, pp. 102–115. DOI: 10.1016/j.gloenvcha.2016.02.003.
 - [149] X. Wen, J. G. Contreras, I. Stadelmann-Steffen, J.-P. Sasse and E. Trutnevyte (2025). High sensitivity to methodological choices when integrating social acceptance data in electricity system modeling. *Applied Energy* **402**, p. 126893. DOI: 10.1016/j.apenergy.2025.126893.
 - [150] J. F. DeCarolís (2011). Using modeling to generate alternatives (MGA) to expand our thinking on energy futures. *Energy Economics* **33** (2), pp. 145–152. DOI: 10.1016/j.eneco.2010.05.002.
 - [151] P. B. Berntsen and E. Trutnevyte (2017). Ensuring diversity of national energy scenarios: bottom-up energy system model with modeling to generate alternatives. *Energy* **126**, pp. 886–898. DOI: 10.1016/j.energy.2017.03.043.
 - [152] F. Neumann and T. Brown (2021). The near-optimal feasible space of a renewable power system model. *Electric Power Systems Research* **190**, p. 106690. DOI: 10.1016/j.epsr.2020.106690.
 - [153] P. Hansen, X. Liu and G. M. Morrison (2019). Agent-based modelling and socio-technical energy transitions: a systematic literature review. *Energy Research & Social Science* **49**, pp. 41–52. DOI: 10.1016/j.erss.2018.10.021.
 - [154] J. Castro, S. Drews, F. Exadaktylos, J. Foramitti, F. Klein, T. Konc, I. Savin and J. van den Bergh (2020). A review of agent-based modeling of climate-energy policy. *Wiley Interdisciplinary Reviews: Climate Change* **11** (4). DOI: 10.1002/wcc.647.
 - [155] W. McDowall and F. W. Geels (2017). Ten challenges for computer models in transitions research: commentary on holtz et al. *Environmental Innovation and Societal Transitions* **22**, pp. 41–49. DOI: 10.1016/j.eist.2016.07.001.

- [156] A. Fossas-Tenas, X. Wen and E. Trutnevyte (2026). Empirical evidence for slack in modeling to generate alternatives. *Energy* **360**, p. 141552. DOI: 10.1016/j.energy.2026.141552.
- [157] M. A. E. van Sluiseveld, J. H. M. Harmsen, N. Bauer, D. L. McCollum, K. Riahi, M. Tavoni, D. P. van Vuuren, C. Wilson and B. van der Zwaan (Nov. 2015). Comparing future patterns of energy system change in 2°C scenarios with historically observed rates of change. *Global Environmental Change* **35**, pp. 436–449. DOI: 10.1016/j.gloenvcha.2015.09.019.
- [158] V. Vinichenko, A. Cherp and J. Jewell (2021). Historical precedents and feasibility of rapid coal and gas decline required for the 1.5°C target. *One Earth* **4** (10), pp. 1477–1490. DOI: 10.1016/j.oneear.2021.09.012.
- [159] S. Link, L. Schneider, A. Stephan, L. Weymann and P. Plötz (2025). Feasibility of meeting future battery demand via domestic cell production in Europe. *Nature Energy* **10**, pp. 526–534. DOI: 10.1038/s41560-025-01722-y.
- [160] K. Zhao, M. A. Wulder, T. Hu, R. Bright, Q. Wu, H. Qin, Y. Li, E. Toman, B. Mallick, X. Zhang and M. Brown (2019). Detecting change-point, trend, and seasonality in satellite time series data to track abrupt changes and nonlinear dynamics: a Bayesian ensemble algorithm. *Remote Sensing of Environment* **232**, p. 111181. DOI: 10.1016/j.rse.2019.04.034.
- [161] D. Sahal (1979). A theory of progress functions. *AIIE Transactions* **11** (1), pp. 23–29. DOI: 10.1080/05695557908974396.
- [162] P. Meyer (1994). Bi-logistic growth. *Technological Forecasting and Social Change* **47** (1), pp. 89–102. DOI: 10.1016/0040-1625(94)90042-6.
- [163] N. Meinshausen (2006). Quantile regression forests. *Journal of Machine Learning Research* **7**, pp. 983–999. URL: <https://www.jmlr.org/papers/volume7/meinshausen06a/meinshausen06a.pdf>.
- [164] E. Byers, V. Krey, E. Kriegler, K. Riahi, R. Schaeffer, J. Kikstra, R. Lamboll, Z. Nicholls, M. Sandstad, C. Smith, K. van der Wijst, F. Lecocq, J. Portugal-Pereira, Y. Saheb, A. Stromann, H. Winkler, C. Auer, E. Brutschin, C. Lepault, E. Müller-Casseres, M. Gidden, D. Huppmann, P. Kolp, G. Marangoni, M. Werning, K. Calvin, C. Guivarch, T. Hasegawa, G. Peters, J. Steinberger, M. Tavoni, D. van Vuuren, A. Al-Khourdajie, P. Forster, J. Lewis, M. Meinshausen, J. Rogelj, B. Samset and R. Skeie (2022). AR6 Scenarios Database. DOI: 10.5281/zenodo.5886912.

-
- [165] IEA (2024). *World energy outlook 2024*. Tech. rep. Paris; France: IEA. URL: <https://iea.blob.core.windows.net/assets/140a0470-5b90-4922-a0e9-838b3ac6918c/WorldEnergyOutlook2024.pdf> (Retrieved: 2025-02-11).
- [166] A. Cheung (Jan. 2025). *Five energy transition lessons for 2025*. URL: <https://about.bnef.com/blog/five-energy-transition-lessons-for-2025/>.
- [167] G. Luderer, S. Madeddu, L. Merfort, F. Ueckerdt, M. Pehl, R. Pietzcker, M. Rottoli, F. Schreyer, N. Bauer, L. Baumstark, C. Bertram, A. Dirnaichner, F. Humpenöder, A. Levesque, A. Popp, R. Rodrigues, J. Strefler and E. Kriegler (2022). Impact of declining renewable energy costs on electrification in low-emission scenarios. *Nature Energy* **7** (1), pp. 32–42. DOI: 10.1038/s41560-021-00937-z.
- [168] D. Bogdanov, M. Ram, A. Aghahosseini, A. Gulagi, A. S. Oyewo, M. Child, U. Caldera, K. Sadovskaia, J. Farfan, L. D. S. N. S. Barbosa, M. Fasihi, S. Khalili, T. Traber and C. Breyer (2021). Low-cost renewable electricity as the key driver of the global energy transition towards sustainability. *Energy* **227**, p. 120467. DOI: 10.1016/j.energy.2021.120467.
- [169] D. Süsser, A. Ceglaz, H. Gaschnig, V. Stavrakas, A. Flamos, G. Giannakidis and J. Lilliestam (2021). Model-based policymaking or policy-based modelling? how energy models and energy policy interact. *Energy Research & Social Science* **75**, p. 101984. DOI: 10.1016/j.erss.2021.101984.
- [170] T. Kazlou, L. Nacke, A. Cherp, A. Pavlenko and J. Jewell (2025). ‘Assessing the feasible expansion of offshore wind in Europe through diffusion heterogeneity’. In: *IST’25*. Lisbon, Portugal.
- [171] J. Gordon, T. Kazlou, J. Jewell and A. Cherp (2026). First-of-a-kind deployment prospects of small modular reactor projects. *In review*.
- [172] G. F. Nemet (2009). Demand-pull, technology-push, and government-led incentives for non-incremental technical change. *Research Policy* **38** (5), pp. 700–709. DOI: 10.1016/j.respol.2009.01.004.
- [173] W. M. Cohen and D. A. Levinthal (1990). Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly*. Special Issue: Technology, Organizations, and Innovation **35** (1), pp. 128–152. DOI: 10.2307/2393553.
- [174] United Nations (2nd Apr. 1992). *Rio declaration on environment and development*. Report. New York: United Nations, Preparatory Committee for the United Nations Conference on Environment and Development, Working Group III.

- [175] F. Denton, K. Halsnæs, K. Akimoto, S. Burch, C. D. Morejon, F. Farias, J. Jupesta, A. Shareef, P. Schweizer-Ries and F. Teng (2022). ‘Accelerating the transition in the context of sustainable development’. In: *Climate change 2022: mitigation of climate change, the Working Group III contribution*, pp. 2816–2915.
- [176] R. J. T. Klein, E. L. F. Schipper and S. Dessai (2005). Integrating mitigation and adaptation into climate and development policy. Three research questions. *Environmental Science & Policy* **8** (6), pp. 579–588. DOI: 10.1016/j.envsci.2005.06.010.
- [177] OECD (2026). *A historic decline in foreign aid: preliminary 2025 ODA data*. Report. Paris: Organisation for Economic Co-operation and Development.
- [178] A. Bhattacharya, V. Songwe, E. Soubeyran and N. Stern (2024). *Raising ambition and accelerating delivery of climate finance. Third report of the Independent High-Level Expert Group on Climate Finance*. Report. London: Grantham Research Institute on Climate Change, the Environment, London School of Economics and Political Science.
- [179] J. Hoffman, C. Born, A. Mahgerefteh, H. Rahman, Z. Greindl, T. Smith, M. Romani, K. Stodulka, A. Bhattacharya and M. Uy (2024). *The state of delivery. Post-COP29 progress report of the global climate finance agenda*. Report. London: Grantham Research Institute on Climate Change, the Environment, London School of Economics and Political Science.
- [180] UNDP (2025). *Human development report 2025. A matter of choice: people and possibilities in the age of AI*. Report. New York: United Nations Development Programme.
- [181] A. Sayer (2019). *Method in social science*. 2nd, revised. London: Routledge. ISBN: 9780415581592. DOI: 10.4324/9780203850374.
- [182] T. Lawson (1998). *Economics and reality*. Routledge. URL: <https://www.routledge.com/Economics-and-Reality/Lawson/p/book/9780415154215> (Retrieved: 2026-09-15).
- [183] G. Iyer, N. Hultman, J. Eom, H. McJeon, P. Patel and L. Clarke (Jan. 2015). Diffusion of low-carbon technologies and the feasibility of long-term climate targets. *Technological Forecasting and Social Change* **90**, pp. 103–118. DOI: 10.1016/j.techfore.2013.08.025.
- [184] T. Napp, D. Bernie, R. Thomas, J. Lowe, A. Hawkes and A. Gambhir (2017). Exploring the feasibility of low-carbon scenarios using historical energy transitions analysis. *Energies* **10** (1), p. 116. DOI: 10.3390/en10010116.

-
- [185] R. Bhaskar (1975). *A realist theory of science*. Routledge. ISBN: 020309073X.
 - [186] B. Latour (1988). *Science in action*. Harvard University Press. URL: <https://www.hup.harvard.edu/books/9780674792913> (Retrieved: 2026-09-15).
 - [187] M. Callon (1986). 'The sociology of an actor-network: the case of the electric vehicle'. In: *Mapping the dynamics of science and technology: sociology of science in the real world*. Ed. by M. Callon, J. Law and A. Rip. London: Macmillan, pp. 19–34. DOI: 10.1007/978-1-349-07408-2_2.
 - [188] E. Trutnevyte (Jan. 2016). Does cost optimization approximate the real-world energy transition? *Energy* **106**, pp. 182–193. DOI: 10.1016/j.energy.2016.03.038.
 - [189] E. Thompson (2022). *Escape from model land*. London: Basic Books. URL: <https://basicbooks.uk/titles/erica-thompson/escape-from-model-land/9781529364873/> (Retrieved: 2026-09-15).
 - [190] T. Jamasb (2007). Technical change theory and learning curves: patterns of progress in electricity generation technologies. *The Energy Journal* **28** (3), pp. 51–72. DOI: 10.5547/issn0195-6574-ej-vol28-no3-4.
 - [191] W. Ye, K. An, S. Zhang, Z. Zhen, W. Cai and C. Wang (2026). Probabilistic analysis of tripling global renewables based on Bayesian-inferred growth dynamics model. *Renewable Energy* **256**, p. 124496. DOI: 10.1016/j.renene.2025.124496.
 - [192] D. Rosenbloom (2020). Engaging with multi-system interactions in sustainability transitions: a comment on the transitions research agenda. *Environmental Innovation and Societal Transitions* **34**, pp. 336–340. DOI: 10.1016/j.eist.2019.10.003.
 - [193] S. J. Davis, N. S. Lewis, M. Shaner, S. Aggarwal, D. Arent, I. L. Azevedo, S. M. Benson, T. Bradley, J. Brouwer, Y.-M. Chiang, C. T. M. Clack, A. Cohen, S. Doig, J. Edmonds, P. Fennell, C. B. Field, B. Hannegan, B.-M. Hodge, M. I. Hoffert, E. Ingersoll, P. Jaramillo, K. S. Lackner, K. J. Mach, M. Mastrandrea, J. Ogden, P. F. Peterson, D. L. Sanchez, D. Sperling, J. Stagner, J. E. Trancik, C.-J. Yang and K. Caldeira (2018). Net-zero emissions energy systems. *Science* **360** (6396). DOI: 10.1126/science.aas9793.
 - [194] J.-F. Mercure (2015). An age structured demographic theory of technological change. *Journal of Evolutionary Economics* **25** (4), pp. 787–820. DOI: 10.1007/s00191-015-0413-9.

- [195] B. K. Sovacool, F. W. Geels, A. D. Andersen, M. Grubb, A. J. Jordan, F. Kern, P. Kivimaa, M. Lockwood, J. Markard, J. Meadowcroft, J. Meckling, B. Moore, R. Raven, K. S. Rogge, D. Rosenbloom, T. S. Schmidt, J. Schot, D. Sharp, J. Stephenson, I. Vormedal and K. Yang (2025). The acceleration of low-carbon transitions: insights, concepts, challenges, and new directions for research. *Energy Research & Social Science* **121**, p. 103948. DOI: 10.1016/j.erss.2025.103948.
- [196] B. A. Sandén and K. M. Hillman (2011). A framework for analysis of multi-mode interaction among technologies with examples from the history of alternative transport fuels in Sweden. *Research Policy* **40** (3), pp. 403–414. DOI: 10.1016/j.respol.2010.12.005.
- [197] L. Nacke, A. Cherp and J. Jewell (2022). Phases of fossil fuel decline: diagnostic framework for policy sequencing and feasible transition pathways in resource dependent regions. *Oxford Open Energy* **1**. DOI: 10.1093/ooenergy/oiac002.
- [198] D. Rosenbloom and A. Rinscheid (2020). Deliberate decline: an emerging frontier for the study and practice of decarbonization. *Wiley Interdisciplinary Reviews: Climate Change* **11** (6). DOI: 10.1002/wcc.669.
- [199] A. Rinscheid, D. Rosenbloom, J. Markard and B. Turnheim (2021). From terminating to transforming: the role of phase-out in sustainability transitions. *Environmental Innovation and Societal Transitions*. DOI: 10.1016/j.eist.2021.10.019.
- [200] A. Sen (1999). *Development as freedom*. Oxford University Press. URL: <https://ideas.repec.org/b/oxp/obooks/9780192893307.html> (Retrieved: 2026-09-15).