



# Dynamic models for clean energy technology diffusion and improvement

**C. Lennart Baumgärtner**

Institute of New Economic Thinking  
Smith School of Enterprise and the Environment  
Worcester College

*Supervisors: Prof. Cameron Hepburn and Dr. Matthew C. Ives*

**Thesis for the degree of Doctor of Philosophy**

School of Geography and the Environment  
University of Oxford  
Oxford, United Kingdom  
August 12, 2025

## Acknowledgements

I would like to thank my supervisors, Prof. Cameron Hepburn and Dr. Matthew C. Ives, as well as Prof. J. Doyne Farmer, for their continuing support. I am very grateful not only for the academic supervision but also for your guidance in more general questions about my career and future. Your support has been invaluable to this thesis and my personal development.

I also thank the different funders sponsoring the presented research, the GEOCEP project under the Marie Skłodowska-Curie grant No. 870245, the Economics of Energy Innovation and System Transitions project, and the Policy Tools for Fast, Effective, Just, and Profitable, Clean Energy Transition project. On top of that, I thank Macrocism Inc. for their generous support.

The presented work has also benefitted hugely from the intense collaboration with my co-authors J. Doyne Farmer, Matthew C. Ives, Rupert Way, Laurin Köhler-Schindler, and Jacquelyn Pless. I could not have asked for a better team, and I am grateful for the inspiration, challenge, and expertise you have provided. This team also includes the spectacular INET administrative staff, particularly Dorothy Nicholas, without whom I would be lost deep within the Oxford labyrinth.

Lastly, I could not have completed this work without the incredible support from my parents, Irmela and Andreas Baumgärtner, my brother Quentin Baumgärtner, and Madalina Graure, who has followed all the ups and downs of my research.

## Abstract

Mitigating global warming is one of the most significant structural transformations undertaken by society. The energy sector is central to this endeavour since it currently accounts for over 73% of global greenhouse gas emissions and trillions of USD in fossil fuel assets. Increased adoption of clean energy technologies, such as solar photovoltaics, onshore wind, and battery-electric vehicles plays an central role in the energy transition by displacing incumbent fossil fuel technologies.

In this thesis, I study dynamic models of future clean energy technology diffusion and cost in the face of uncertainty. I answer key questions about the energy transition whilst advancing statistical models of technological change.

In my first paper, I critically examine recent models which have forecasted limits to the diffusion of solar photovoltaics and onshore wind, finding that many results are not robust due to underlying uncertainty and the lack of statistical rigour in the assessments. Global solar and wind deployment have not significantly departed from their exponential growth trajectory so far, and there are no reliable signs when they will do so.

For my second paper, I apply more reliable statistical methods to develop forecasts for national solar and wind costs. Based on their decomposition, solar costs will likely decline further, but wind may soon reach its floor cost. Solar levelised costs of electricity in different countries will most likely converge in the future due to a high degree of correlation between countries. Wind costs, however, will likely continue to display cross-sectional variations of almost an order of magnitude due to differences in natural resources.

My third paper uses similar forecasts to develop dynamic investment strategies in technology portfolios under uncertainty and learning. For instance, in the case of battery-electric passenger vehicles, assuming low discount rates and negligible switching costs, we are now at a tipping point where it is no longer cost-optimal to invest in traditional combustion-based vehicles. However, with an discount rates above 2%, continued policy support is required if the transport sector is to reach this tipping point.

# Table of Contents

<b>I</b>	<b>Introduction</b>	<b>7</b>
<b>II</b>	<b>Literature review</b>	<b>37</b>
II.1	Technological change in the energy transition . . . . .	37
II.1.1	Deployment of clean energy technologies . . . . .	38
II.1.2	Cost of clean energy technologies . . . . .	40
II.1.3	Integrated Cost and Deployment models . . . . .	42
II.2	Empirical analysis of technological change . . . . .	43
II.2.1	S-curve models of technology deployment . . . . .	44
II.2.2	Experience curve models of technology cost . . . . .	45
II.2.3	Investment allocation models . . . . .	48
II.2.4	Validation of technological change models . . . . .	52
<b>III</b>	<b>Paper 1: The need for better statistical testing in data-driven energy technology modeling</b>	<b>85</b>
III.1	Introduction . . . . .	85
III.2	Examples of overfitting and selection bias . . . . .	89
III.3	Testing an energy technology diffusion model from the recent literature . .	93
III.4	Key Takeaways: Making reliable statistical inferences . . . . .	100
<b>IV</b>	<b>Paper 2: Will national renewable costs continue declining?</b>	<b>111</b>
IV.1	Introduction . . . . .	112
IV.2	Decomposition of the LCOE . . . . .	113
IV.3	Forecasting the components of LCOE . . . . .	118
IV.4	Forecasting national costs . . . . .	124
IV.5	Discussion . . . . .	128
IV.6	Methods . . . . .	130

IV.6.1	Data collection and cleaning . . . . .	130
IV.6.2	LCOE decomposition . . . . .	132
IV.6.3	Solar module and wind turbine component model . . . . .	133
IV.6.4	Balance of system component model . . . . .	134
IV.6.5	Capacity factor component model . . . . .	134
IV.6.6	LCOE forecasts . . . . .	135
IV.6.7	IAM scenario comparison . . . . .	136

## V Paper 3: Innovation Bandits: A Dynamic Portfolio Strategy with

<b>Endogenous Rewards</b>		<b>145</b>
V.1	Introduction . . . . .	146
V.1.1	Related Portfolio Optimization Literature . . . . .	151
V.1.2	Summary of Contributions . . . . .	154
V.1.3	Outline . . . . .	156
V.2	A portfolio optimization problem with endogenous rewards . . . . .	157
V.2.1	Technology Cost Evolution . . . . .	157
V.2.2	Objective function . . . . .	159
V.3	Index strategies for continuous innovation bandits . . . . .	161
V.3.1	The Gittins Index . . . . .	161
V.3.2	Gittins Indices for Innovation Bandits . . . . .	163
V.3.3	Numerical Solutions to the Gittins Index Strategy . . . . .	164
V.3.4	Gittins Index vs. Expected Net Present Cost . . . . .	167
V.3.5	Gittins Index for Mean-Reverting Bandits . . . . .	170
V.4	Application to Light Vehicle Investments . . . . .	171
V.4.1	Light Vehicle Experience Curves . . . . .	172
V.4.2	Light Vehicle Investment Strategy . . . . .	174
V.4.3	Dynamic Policy Choices . . . . .	177
V.5	Discussion . . . . .	179
V.5.1	Dynamically Optimal Innovation Portfolios . . . . .	179

V.5.2	Public Policy Implications . . . . .	180
V.5.3	Opportunities for Future Work . . . . .	181
V.6	Conclusion . . . . .	184
<b>VI</b>	<b>Conclusion</b>	<b>197</b>
<b>A</b>	<b>Supplementary Materials for</b> <i>The need for better statistical testing in data-driven energy technology modeling</i>	<b>215</b>
A.1	Supplementary discussion on the examined statistical model selected from recent literature . . . . .	215
A.1.1	Detailed explanation of the model . . . . .	215
A.1.2	The reliability of inflection point estimates . . . . .	216
A.1.3	The application of the examined statistical model to other technologies . . . . .	218
A.2	Additional Figures . . . . .	223
<b>B</b>	<b>Supplementary Materials for</b> <i>Will national renewable costs continue declining?</i>	<b>225</b>
B.1	Notation overview . . . . .	225
B.2	Description of national data collection and cleaning . . . . .	226
B.2.1	Capacity and generation data . . . . .	226
B.2.2	LCOE data . . . . .	226
B.2.3	Total investment costs . . . . .	230
B.2.4	Module and Turbine costs . . . . .	230
B.2.5	Balance of system cost . . . . .	231
B.2.6	Capacity factor . . . . .	231
B.2.7	Operating and maintenance costs . . . . .	232
B.2.8	Capital recovery factor . . . . .	233
B.3	LCOE model development . . . . .	234
B.3.1	LCOE decomposition and aggregation . . . . .	235

B.3.2	Module cost forecasts . . . . .	242
B.3.3	Turbine cost forecasts . . . . .	243
B.3.4	Solar BOS cost forecasts . . . . .	247
B.3.5	Wind BOS cost forecasts . . . . .	260
B.3.6	Investment cost forecasts . . . . .	266
B.3.7	Capacity Factor forecasts . . . . .	269
B.3.8	Operating & Maintenance cost forecasts . . . . .	275
B.3.9	Additional Cross-checks . . . . .	278
B.4	IAM scenario comparison . . . . .	282
B.4.1	Scenario data collection . . . . .	283
B.4.2	Initial value correction . . . . .	283
B.5	Estimated model parameters . . . . .	285
B.6	Additional Figures . . . . .	285

**C Supplementary Materials for *Innovation Bandits: A Dynamic Portfolio Strategy***

	<i>with Endogenous Rewards</i>	<b>296</b>
C.1	Model assumptions . . . . .	296
C.2	Proofs: Gittins indices . . . . .	297
C.2.1	Mean-reverting bandits . . . . .	298
C.2.2	Innovation bandits . . . . .	300

# Chapter I

## Introduction

In 2024, the global average temperatures relative to the 1850-1900 baseline exceeded the 1.5°C threshold for the first time, making it the hottest year on record ([World Meteorological Organization, 2025](#)). While annual temperature fluctuations are expected, Earth is at high risk of breaching the long-term 1.5°C global warming ambition set by the Paris Agreement ([Bevacqua et al., 2025](#); [Cannon, 2025](#)). Rising global temperatures pose substantial societal and ecological risks, including extreme weather events and diminishing natural habitats ([IPCC, 2023](#)). Reducing greenhouse gas emissions across regions and sectors is an important lever to mitigate these risks ([IPCC, 2022](#)).

Beyond reducing and shifting demand for greenhouse gas-intensive products and services, *technological change*<sup>1</sup> lies at the heart of most Net Zero mitigation pathways ([IPCC, 2022](#)). In these pathways, high-emission technologies are displaced by clean technologies that do not emit greenhouse gases, and innovations capable of achieving net negative emissions are introduced.

Technological change is particularly important for energy sector transformation. In 2021, the energy sector accounted for over 73% of global emissions, primarily from electricity and heat generation (~ 33%), transportation (~ 17%), and manufacturing use (~ 13%) ([Ritchie et al., 2023](#)). These emissions largely stem from fossil-fuel combustion, such as coal and gas electricity, petrol-based transportation, or the natural gas-based heating of buildings and industrial processes. Transitioning to a clean energy system requires a combination of increased efficiency, demand management, and the replacement of fossil-fuel-based technologies with clean alternatives. For example, clean renewable electricity sources (solar photovoltaics, on/offshore wind, nuclear, etc.) can reduce emissions by

---

<sup>1</sup>Following [Arthur \(2011\)](#), we define a technology as a device or process that uses resources to provide a service to its users. Technologies may further cause undesired secondary effects, such as greenhouse gas emissions and other environmental damages.

replacing some of our fossil-fuel-based sources (coal electricity, gas turbines, etc.)

The economic cost of such technology replacements is a key determinant of the scale and speed of the energy transition (Stern, 2014). Alongside other factors, such as land and material availability, geopolitics, or social acceptance, economic cost confines the energy transition. One way to see this is the striking correlation between personal energy use and GDP per capita (OWID, 2023). For every doubling of energy use, GDP per capita approximately doubles, independent of a country’s total GDP. While some countries have seen economic growth without a strong increase in energy use (Ritchie, 2021), it is clear that access to affordable energy is a central driver of global inequality. The same is true for clean energy sources.

On the surface, the required scale, speed, and cost of clean energy technology deployment are daunting. In electricity alone, achieving Net Zero requires an additional 16.7 PWh/year of clean electricity generation—equivalent to 61% of the 27.5 PWh global electricity demand in 2023 (Ember, 2024). At current deployment rates of 404 TWh per year (IEA, 2024a), replacing fossil fuels with renewables would take until 2064, far beyond the Net Zero timeline of a 2°C warming target (Byers et al., 2022). Moreover, at present costs of approximately 500 USD per MWh of annual electricity production (IRENA, 2021), the energy transition would require over 8 trillion USD in investment, excluding the cost of grid infrastructure. Furthermore, given the typical 40-year asset lifetime (Pfeiffer et al., 2018), over 65 trillion USD worth of existing fossil fuel assets would need to be decommissioned or retrofitted.

However, it is important to put these numbers into context. Firstly, we spend around 1.1 trillion USD annually to maintain the current fossil-fuel-based energy system, based on the average over the last ten years (IEA, 2024b). This is in addition to the higher operational costs and lower fuel efficiency related to fossil-fuel-based energy. For example, petrol-based transportation is around 25% efficient (final to useful energy efficiency), while electricity-based transportation has an efficiency of approximately 80% (Way et al., 2022).

Past progress in clean energy technology deployment and cost offers a further source

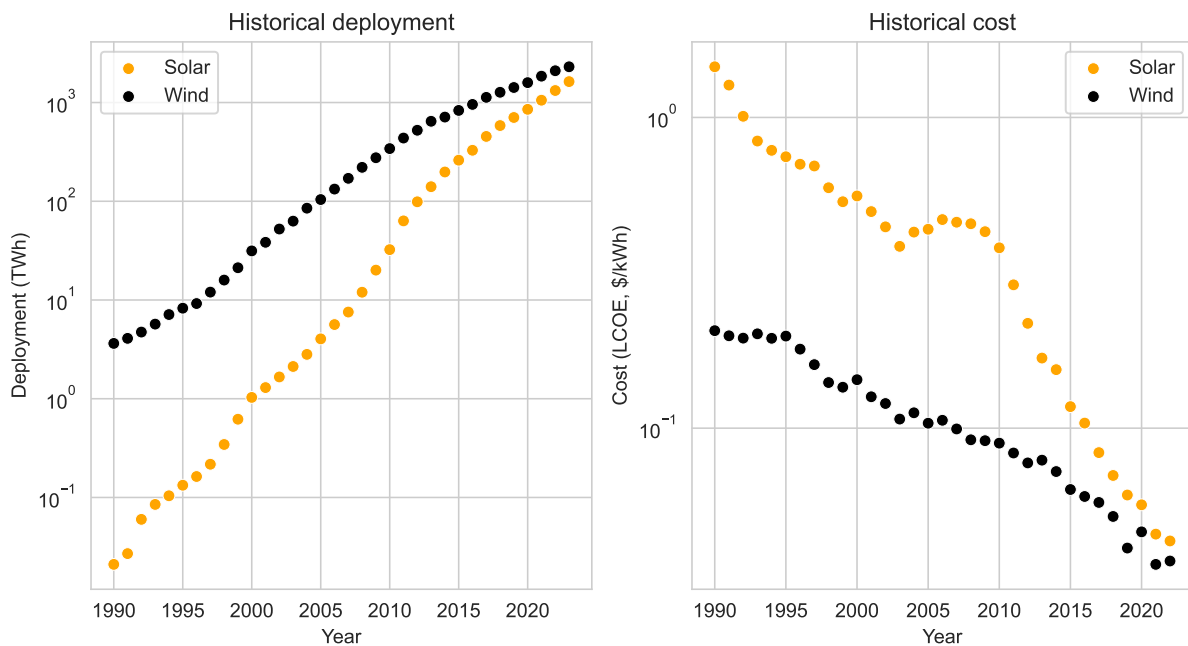


Figure I.1: **Historical deployment growth and cost decline of solar photovoltaic and onshore wind.** Since 1990, electricity production from solar photovoltaics has increased by five orders of magnitude and three orders of magnitude from onshore wind. Over the same period, solar costs have decreased by two orders of magnitude and wind by one. Data from [Way et al. \(2022\)](#).

of optimism. Figure I.1 shows the historical deployment growth and cost decline of solar photovoltaics and onshore wind. Over the past three decades, solar electricity generation has increased by five orders of magnitude, while costs have fallen by two orders of magnitude. Wind power has also expanded significantly, with three orders of magnitude in deployment growth and one order of magnitude of cost reduction ([Way et al., 2022](#)). The logarithmic scale in Figure I.1 indicates that both deployment growth and cost decline have been approximately exponential, with roughly 30% year-on-year growth in solar and 20% in wind. Under such an exponential model, the rapid energy transition appears achievable. If these trends continue, solar alone could supply the required 16.7 PWh/year by 2035 at a fraction of today’s costs ([Way et al., 2022](#); [Farmer and Makhijani, 2010](#)).

In reality, the future trajectory of renewable energy deployment and costs will deviate from a simple exponential model. Constraints such as supply chain bottlenecks ([van de Ven et al., 2021](#)), land-use limitations ([Olivetti et al., 2017](#)), and social acceptance ([Ellis et al., 2023](#)) may slow progress. Similarly, clean energy costs may not continue to follow

an exponential trend. Floor costs from material constraints appear likely (Hsieh et al., 2019), and different technology components may not be subject to learning at all (Ferioli et al., 2009). These factors introduce uncertainty into the pace and economic feasibility of the energy transition. One research direction to reduce and quantify these uncertainties is to study technological change in the energy transition.

The dual objective of this thesis is thus to investigate the reality of future deployment and cost of clean energy technologies while making methodological advances to dynamic models used in the field. Specifically, the focus is on the growth and cost trajectories of individual clean energy technologies and the corresponding dynamic modeling approaches. This emphasis prioritizes individual technology adoption, rather than analysing overarching energy system costs, which inherently depend on interactions among multiple technologies and broader societal dynamics. As discussed later, this focus on individual technologies limits the scope of the results by excluding a full system perspective. Nevertheless, the models and insights developed here provide important inputs to larger system models and offer a basis for validating system-level findings.

Chapter II provides an overview of the current literature on technological change in the energy transition and associated dynamic models. Across the three papers presented in Chapters III, IV and V, I discuss the deployment and cost of utility-scale solar photovoltaics, onshore wind, and battery-electric vehicles. These three technologies are particularly relevant for the energy transition due to their large potential for deployment growth and cost decline, although the energy transition will also require other clean technologies, such as grid batteries, nuclear power, or electrolyzers (IPCC, 2022).

The methods used and developed throughout this thesis focus on modelling the dynamic and stochastic aspects of technological change. The validation and application of the models developed in this thesis are not limited to the energy transition. In fact, assessing how we can apply findings from previous technological transitions to the energy transition is a central theme of this thesis.

## Models of deployment and cost in technology transitions

Understanding when and how technology transitions will unfold and how we may change their trajectory is a fascinating research topic in its own right. These dynamics shed light on how our society is progressing and what the emerging dynamics of this complex system are (Farmer, 2024).

The main approaches for quantifying technological change are expert surveys, mechanistic models, and stochastic models based on historical data (Meng et al., 2021; McNerney et al., 2011; Gambhir et al., 2021).

- *Expert-based methods* rely on structured quantitative and qualitative surveys, such as the Delphi method (Rowe and Wright, 1999). These methods are particularly useful when historical data is sparse or unavailable. They also frequently inform the other two methods through qualitative insights.
- *Mechanistic models* take a bottom-up approach by decomposing technological change (e.g., cost reductions or adoption rates) into its distinguishing features. These models explore how shifts in these features translate into technological trends. The underlying assumptions can be based on expert input, historical data, or theoretical arguments (Kavlak et al., 2018; McNerney et al., 2011). Mechanistic models help quantify the interaction of different technological change features and their contribution to the long-term trend. They also oftentimes provide more insights into the causal drivers of technological change than their model alternatives and allow for subsequent 'what-if' analyses (Nemet, 2006).
- *Statistical models* or *data-based models* analyse historical data to identify patterns and trends that can be extrapolated into the future (Clements and Hendry, eds, 2008). Despite not explaining the root causes of technological change, they have been shown to provide more accurate forecasts of future changes than model alternatives (Meng et al., 2021; Nemet et al., 2017). Perhaps most importantly, statistical models

can be tested to assess their accuracy and reliability (Farmer and Lafond, 2016).

For this thesis, I focus on statistical models. Statistical models are well-suited to study clearly defined properties of technological change for which sufficient historical data is available. In these settings, they outperform their alternatives and provide both more accurate and more reliable results (Meng et al., 2021; Nemet et al., 2017; Nagy et al., 2013; Clements and Hendry, eds, 2008). Conversely, statistical models are ill-suited to answer research questions without clear definitions of what is being studied and insufficient data since they cannot be calibrated precisely. In our case, the future deployment and cost of solar photovoltaics, onshore wind, and battery-electric vehicles are well-defined, and we have multiple decades of historical data available. Thus, statistical models are an appropriate choice for studying future trends in these technologies. For the remainder of this thesis, I, therefore, concentrate on statistical models and refer readers interested in expert elicitation and mechanistic models to Rowe and Wright (1999); Neij (2008); Sabatier (1986); Nemet (2006); Funk and Magee (2015).

Despite their advantages, statistical models are not without limitations. Large parts of this thesis are therefore dedicated to the empirical and theoretical limitations of our findings. Stochastic models can be subject to various methodological pitfalls, such as biases, over- and under-fitting, or the use of ad-hoc assumptions (Chapters III and IV). This means that, in many cases, there is not enough data to make statistically significant statements. This is particularly evident in causal relationships between drivers of technological change (Chapter IV). As a result, model choices are usually not unique or require additional assumptions (Chapter V). As is often the case in social science, there is no perfect model choice.

Within stochastic models, the most relevant models to this thesis are phenomenological descriptions of technological change. *S-curves* and *experience curves* play a central role in describing the deployment and cost of new technologies, respectively.

S-curve phenomena were first discovered in biology, studying the growth and decline of species (Grübler, 1990; Lotka, 1956; Volterra, 1927). S-curves have been shown to

represent technological diffusion well in over 40 examples, ranging from steel furnaces, canals, ATMs, and steamships to mobile phones and internet users (Comin et al., 2006; Wagenvoort et al., 2025). The most commonly used S-curve is the logistic growth model,

$$S = f(t) := \frac{L}{1 + e^{-k(t-t_0)}}, \quad (\text{I.1})$$

where  $S$  represents technology deployment,  $k$  is the initial exponential growth rate,  $L$  is the asymptotic saturation level, and  $t_0$  is a location parameter, also called the inflection point. Initially, the deployment  $f(t \ll t_0)$  is very small but grows approximately exponentially at a year-on-year rate of  $k$ . Eventually, this growth slows down as it reaches its maximum at  $t_0$ . Growth approaches zero as the deployment reaches the asymptote  $L$ . In this model, we can interpret technologies as biological species competing for limited resources (Grübler, 1990; Bhargava, 1989). Of course, this analogy only goes so far since technologies are not only competitive but can also be complementary (Sinsel et al., 2020).

Experience curve phenomena were first discovered in aircraft manufacturing during the Second World War. Wright (1936) discovered how production costs decline as cumulative production increases:

$$C(t) = C_0 \cdot \left( \frac{Z(t)}{Z_0} \right)^\omega, \quad (\text{I.2})$$

where  $C(t)$  is the cost at time  $t$ ,  $Z(t)$  is the cumulative production at time  $t$ , and  $\omega$  is a technology-specific learning rate.  $C_0$  and  $Z_0$  are the initial cost and cumulative production at the start of the experience curve. Under this model, technology costs may be high initially ( $C_0$ ) but can decline rapidly when their market grows. This behavior emerges theoretically from a random search problem, subject to some technology characteristics (Auerswald et al., 2000; McNERNEY et al., 2011). Empirical analyses have confirmed the predictive power of experience curves across more than 50 technologies (Lafond et al., 2018). Of course, technology costs depend on many more factors than their cumulative deployment, such as economies of scale and policy intervention (Kavlak et al., 2018; Nemet, 2006). Nevertheless, experience curves are a good predictor of future costs since, with

sufficient historical data, they are relatively robust to overfitting and selection bias (Nagy et al., 2013).

We can further combine S-curves and experience curves to obtain a more systematic representation of a technology ecosystem and study the interaction of both phenomena (Grübler et al., 1999; Way et al., 2022). This is a challenging task since technology cost and deployment are not independent of each other and other technologies. One of the core issues is causality (Nordhaus, 2014): Does experience stimulate demand, or does demand stimulate experience? The answer lies somewhere in between. Declining costs are partially due to learning over time and an increase in deployment. Vice versa, increasing deployment is partially due to declining costs (Lafond et al., 2020).

As a result, analysing deployment and cost contemporaneously usually requires additional assumptions about decision-making processes. In particular, techno-economic models frequently make assumptions about how decision-makers choose between different technologies that compete or complement each other. The predominant assumption in economics is that agents optimise their decisions to maximise utility, leading to market equilibrium (Nordhaus, 1992; Ramsey, 1928; Ingersoll, 1987):

$$\sup_{(S_i(t))_{i=1,2,\dots}} \sum_{t=0}^{\infty} e^{-\rho t} U[(S_i(t))_i], \quad (\text{I.3})$$

where  $(S_i(t))_{i=1,2,\dots}$  is the deployment of technology  $i$  at year  $t$ ,  $U[\square]$  is a convex utility function, and  $\rho > 0$  is a discount rate representing time-preference. Depending on the choice of the utility function, this supremum has a unique maximum that can be identified computationally.

While the optimality assumption is widely used, it has been critiqued on both theoretical and empirical grounds. Empirically, various studies indicate that decision-makers may not optimise their utility or that utility is heterogeneous without a Pareto-optimal solution (Farmer et al., 2015; Trutnevyte, 2016; Barbrook-Johnson et al., 2024). Theoretically, investment decisions in one technology over another are subject to discontinuous

and counterintuitive behaviour, depending on the utility function (Dietz and Hepburn, 2013; Way et al., 2019). Nevertheless, optimization-based models remain relevant for understanding technology adoption dynamics in business and policy settings, in replicating decision-makers' general inclination towards lower-cost options. It is helpful to understand optimal pathways and solutions, even if these will not ultimately be implemented (Sinha et al., 2024). Furthermore, optimization models are helpful in multi-criteria decision-aid systems to disaggregate consumer and investor preferences (Jacquet-Lagrèze and Siskos, 2001).

This thesis advances S-curve models, experience curve models, and optimization models with a strong focus on real-world applications in the energy transition. Importantly, modelling choices are shaped by both theoretical considerations and practical constraints. While theoretical models ideally capture fundamental system dynamics, they also need to be implemented and calibrated with the available data when used in practice (Meade and Islam, 1998; Petropoulos et al., 2022; Powell, 2011). In other words, a model that works well in theory may not work in practice because it requires too many resources to run the model and cannot be calibrated with existing datasets.

There are a number of fields that have made significant progress by focusing model development on their practical applications. Meteorology, for example, has reached an unprecedented degree of weather forecast accuracy by developing models that work best under existing resource and data constraints while simultaneously lowering these constraints through increased computational power and input data (Alley et al., 2019). One field that has put the application of statistical models at the forefront is that of machine learning (Kelleher et al., 2020). This thesis incorporates some of these newer machine learning techniques to refine models of technological change.

Arguably, the most important technique is that of *model validation*. It takes centre stage in chapters III-V and, as we will see, is the critical determinant in our model choices. Validation is key to assessing the reliability of models (Baumgärtner et al., 2024). This applies to most types of models, such as prediction models or historical ex-post analyses.

Model outputs can be misleading due to statistical issues. For example, overfitting a model means that model results reflect random noise in the underlying data instead of (causal) correlations. While no validation method is perfect, it can identify statistical issues and build more confidence in a model's result.

In the context of statistical models, model validation reduces the risk that the model represents the noise or uncertainty underlying the data rather than the drivers (signal) in our data. Validation exercises generally require assumptions about the stochastic model underlying the noise in our data. This is true of all statistical models, even if the stochastic assumptions are not stated explicitly. For example, modellers distinguish between exogenous noise that is independent of the model variables and endogenous noise that depends on the statistical signal itself (Salo et al., 2024). The goal of model validation is then to identify the type and magnitude of noise and make sure the model's results are robust in the face of the noise.

As we will see, there are two types of model validation. We can validate models empirically by performing statistical tests on model results. We can also validate models theoretically by assessing their underlying assumptions and verifying the internal consistency. Chapters III and IV focus on empirical validation of prediction models for technology deployment and cost. Chapter V explores the theoretical validation of optimization models that simulate optimal deployment decisions under varying cost assumptions. Through these validation efforts, this thesis aims to develop more robust models that contribute to a better understanding of the energy transition.

## **Research objectives: A new narrative for the energy transition**

Model validation exercises have already uncovered several misconceptions about the energy transition. For example, Creutzig et al. (2017) show that past energy system scenarios have consistently underestimated the future growth of solar photovoltaics, while Xiao et al. (2021) demonstrate that similar scenarios have also overestimated the future costs of renewable energy sources. Perhaps most strikingly, Way et al. (2022) reveal that, due to

innovation in clean energy technologies, the long-standing notion that decarbonization will be costly<sup>2</sup> may not hold for the energy sector. Instead, rapid decarbonization has the potential to reduce energy system costs by 12 trillion USD (present value).

These findings, along with recent advances in clean technologies, have shifted the central discussions around technological change in the energy transition. In the past, a lot of research has focused on setting the right technological change targets and determining global warming goals (McLaren and Markusson, 2020). While these questions remain important, researchers and policymakers today engage with a much broader set of issues (IPCC, 2022), even within the narrow scope of technology cost and deployment. For example, a significant amount of work has gone to assess and plan the land and material requirements for scaling up clean technologies (and their potential impact on costs) (de Castro et al., 2013; Olivetti et al., 2017; Sørensen, 1975), investigating the consequences for labour markets and industrial policies (Rodrik, 2014; Mealy and Teytelboym, 2022; Bücker et al., 2025), and studying the broader political implications of the transition (Hallegatte et al., 2023; Nature Energy, ed, 2024; Markard, 2018).

For this thesis, I focus on three topics that serve the dual objective of advancing the literature on stochastic models of technological change and the study of the energy transition. By developing new methods for analyzing technological change, we can provide novel insights about the energy transition. This thesis concentrates on the broad systemic dynamics of the energy transition. As such, I hope the research findings will have broad applications across countries and sectors. There is significant scope for future research to explore the various implications of these findings in specific contexts at the level of individual decision-makers. I will discuss future research in more detail in the concluding chapter of this thesis.

---

<sup>2</sup>For his integrated assessment model *DICE*, Nordhaus (1992) makes the bold assumption that “The major choice faced by the economy in the DICE model is whether to consume goods and services, to invest in productive capital, or to slow climate change.”

## First paper ([Baumgärtner et al., 2024](#))

In “The Need for Better Statistical Testing in Data-Driven Energy Technology Modeling,” Rupert Way, Mathew C. Ives, J. Doyne Farmer, and I examine: *How to assess the feasibility of fast and cheap clean energy technology diffusion, particularly in solar photovoltaics, onshore wind, and electric batteries?* This paper was published in Joule and is presented in Chapter III.

A growing body of research has examined the *feasibility* (or, interchangeably, the plausibility) of fast decarbonization and declining clean technology costs ([Sognaes, 2022](#)). Future values of a technology’s deployment or cost are considered feasible if there is a non-zero probability of those values occurring; that is, feasibility studies aim to assess what future technological change is possible to achieve. Examples include [Odenweller et al. \(2022\)](#), who investigate how quickly green hydrogen production can be ramped up; [Nemet et al. \(2023\)](#), who study the growth ambition of carbon direct removal; and [Cherp et al. \(2021\)](#), who explore viable growth of solar photovoltaics and onshore wind. These studies are important for developing detailed energy transition strategies and ensuring that ambitious plans are still realistic.

However, the results from feasibility assessments face multiple limitations. First, these analyses add to an already complex landscape of “scenarios”, “projections”, and “predictions” of the energy transition. This has created some confusion around their interpretation ([Bunn and Salo, 1993](#); [Sognaes, 2022](#)). Second, feasibility analyses have yielded conflicting results for some technologies. While some studies argue that rapid energy transition through renewables is infeasible ([Smil, 2016](#)), others present opposing conclusions ([Wilson, 2012](#)). This further adds to the perplexity of feasibility spaces in technological change and raises the question of why these results differ and how to assess feasibility with statistical rigour.

The key challenge in answering this question lies in the multitude of factors driving the deployment and cost of clean technologies, such as manufacturing scale, labour markets,

and policy support. On top of that, solar photovoltaics, onshore wind, and batteries all exhibit non-linear and non-marginal changes in their cost (experience curves) and deployment (S-curves). These changes are generally difficult to predict (Petropoulos et al., 2022). Most techno-economic methods, such as general equilibrium modelling, struggle with these dynamics, as they are ill-suited for handling non-linear feedback effects (Farmer et al., 2015). Instead, historical patterns of technological diffusion and cost reductions can provide insights into how novel technologies may evolve (Grübler, 1990; Grübler et al., 1999). However, the limited availability of historical data, combined with the non-linearity of these models, necessitates rigorous statistical testing to improve the reliability of their conclusions.

In this paper, I highlight multiple examples from the recent energy technology modelling literature where conclusions are unreliable due to insufficient data and a lack of statistical validation. For instance, claims that solar photovoltaics and onshore wind deployment are unlikely to continue accelerating and meet international growth targets are not supported by robust empirical evidence. Instead, the data shows that recent wind and solar deployments remain statistically consistent with their historically exponential growth trajectories.

To avoid such pitfalls and spurious conclusions, we emphasize the importance of out-of-sample testing, as it minimizes the reliance on strong assumptions about the underlying statistical properties of the data.

## Second paper

In “Will national renewable costs continue declining?,” J. Doyne Farmer and I apply these principles of robust model validation to answer the question: *What are the future national costs for solar photovoltaic and onshore wind?* This paper is included in Chapter IV.

Beyond feasibility, the actionability and equity implications of decarbonization have become central concerns. Even if energy decarbonisation has positive GDP impacts, it is unclear how these are distributed. One area of focus is the variation of decarbonization costs and pathways across countries (Plazas-Niño et al., 2022; Mulugetta et al., 2022;

[Mercure et al., 2021](#)). Renewable energy costs differ by up to an order of magnitude between countries ([IRENA, 2021](#)). This creates new risks and opportunities to participate in the emerging green supply chain ([Bowen and Hepburn, 2014](#); [Mealy and Teytelboym, 2022](#); [Grubb et al., 2021](#)). For example, energy costs comprise 20-40% of production costs in the steel sector ([World Steel Association, 2021](#)). Future local energy costs will be a key determinant of how fast the steel sector will decarbonise ([Xu et al., 2023](#)), but also where new manufacturing locations will be built. Forecasting future national renewable costs is imperative to assessing national decarbonisation costs and broader industrial strategies.

Through developing forecasts for national renewable costs, this paper contributes to the energy transition and technological change literature in three key elements. First, we have compiled a novel dataset on historical solar and wind costs, along with their components, covering over 100 countries. This represents the most comprehensive dataset used in renewable cost forecasting to date. Second, we forecast individual components of the levelized cost of electricity (LCOE) and find that not all components are subject to long-term cost decline. We discover a floor cost for wind investment costs that we are bound to approach soon. Third, we identify that national renewable costs are highly correlated with each other. This means that global learning effects are dominant over local effects and future national costs depend on the global decarbonization pathway.

Our dataset and forecasts are based on decomposing the LCOE into its components. We collect data for each component and make forecasts for each component and country. Each forecast is stochastic in nature. This way, we capture the future expected costs of each component, as well as the uncertainty associated with it. For model validation, it is important that the model dynamics for each country are equivalent. Equivalent model dynamics substantially reduce the risk of overfitting our limited data and make the model more widely applicable.

Our forecasts of national renewable energy costs provide valuable insights for both policymakers and the broader energy modelling community. Anticipating future national clean technology costs is critical for setting effective decarbonization strategies and design-

ing appropriate policy support mechanisms. While some policies, such as international trade regulations, fall outside the scope of this analysis, our model offers an important first step toward understanding the impact of different policy measures. We observe that the current integrated assessment modelling literature (Byers et al., 2022) systematically overestimates future solar costs while underestimating the uncertainty surrounding future renewable energy costs. Our approach provides a viable alternative that can be integrated into energy system scenarios to improve the accuracy and robustness of renewable energy projections.

### Third paper

In “Innovation Bandits: A Dynamic Portfolio Strategy with Endogenous Rewards,” co-authored with Laurin Köhler-Schindler and Jacquelyn Pless, I investigate *How to optimally allocate resources across a portfolio of technologies that evolve according to a stochastic experience curve?* This paper is available in Chapter V.

This third and final research question emerges naturally when stochastic cost forecasting models are applied to multiple competing technologies. Optimal resource allocation amongst technologies is not a new research question in the context of energy transition but has been around since the first international climate conferences and IPCC reports on climate change (Girod et al., 2009). However, due to the high historical costs of clean energy technologies, the associated optimization problem focused mostly on identifying low-cost decarbonisation pathways that required the least amount of policy support (Pindyck, 2013). Today, this narrative is changing towards identifying investment decisions that both reduce carbon and create economic growth opportunities (Mealy and Teytelboym, 2022; Farmer et al., 2015). In this context, portfolio optimization can help identify tipping points in a technology ecosystem where clean technologies are the preferred economic choice (Mealy et al., 2023; Farmer et al., 2019).

A key challenge with portfolio optimization for competing technologies is that future technology costs and their associated uncertainty are endogenous to the portfolio choice

(Way et al., 2019; Cowan, 1991). Endogenous uncertainty, in particular, is a challenge to the validation of optimization models because optimal choices can not be verified ex-post (Salo et al., 2024). This means that traditional optimization techniques, such as prioritizing technologies with the lowest expected net present cost, may no longer be optimal. Since the uncertainty means that future costs are subject to noise, the optimal resource allocation strategy reacts dynamically to the realised technology costs.

In this paper, we develop a novel approach to this dynamic optimization problem, assuming that resource allocations can be adjusted arbitrarily over time. Borrowing a framework from reinforcement learning, the Multi-Armed Bandit framework, we use the Gittins index to identify the most promising technology to invest in at any given moment (Gittins et al., 2011). The Gittins index quantifies the option value of investing in a particular technology. While calculating this index is generally difficult and not always feasible, we demonstrate how it can be computed numerically for technologies evolving under an experience curve or mean-reversion process. While there are many assumptions that we have to take for this model to compute, this work represents, to our knowledge, the first actionable application of Multi-Armed Bandit methods in technology investment studies.

To illustrate our method, we compare different investment strategies for battery-electric vehicles (BEVs) and conventional internal combustion engine vehicles (ICEs). BEV costs evolve under an experience curve from reduced battery costs. ICE costs, on the other hand, are not subject to significant cost decline and can be modelled with a mean-reverting stochastic process. We find that for high discount rates, the Gittins index strategy aligns with the traditional static approach that estimates the expected net present cost of a technology portfolio. However, for low discount rates, the optimal strategy favours investment in BEVs despite their currently higher costs due to the upside potential for future cost reductions (Dixit and Pindyck, 1994).

This insight has direct policy implications. Firstly, it supports policy favouring technologies with large associated uncertainty. Successful technology policies inevitably fail

from time to time ([Rodrik, 2014](#)). Secondly, the current macroeconomic environment with high interest rates discourages long-term investment in emerging technologies. This means that accelerated BEV adoption will likely require continued policy support. However, this support need may subside once BEV costs fall below a critical threshold or macroeconomic conditions change.

## References

- Alley, Richard B., Kerry A. Emanuel, and Fuqing Zhang**, “Advances in Weather Prediction,” *Science*, January 2019, *363* (6425), 342–344.
- Arthur, W. Brian**, *The Nature of Technology: What It Is and How It Evolves*, first trade paperback ed., New York, London, Toronto, Sydney: Free Press, 2011.
- Auerswald, Philip, Stuart Kauffman, José Lobo, and Karl Shell**, “The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning by Doing,” *Journal of Economic Dynamics and Control*, March 2000, *24* (3), 389–450.
- Barbrook-Johnson, Pete, Jean-François Mercure, Simon Sharpe, Cristina Peñasco, Cameron Hepburn, Laura Diaz Anadon, J. Doyne Farmer, and Timothy M. Lenton**, “Economic Modelling Fit for the Demands of Energy Decision Makers,” *Nature Energy*, March 2024, *9* (3), 229–231.
- Baumgärtner, C. Lennart, Rupert Way, Matthew C. Ives, and J. Doyne Farmer**, “The Need for Better Statistical Testing in Data-Driven Energy Technology Modeling,” *Joule*, September 2024, *8* (9), 2453–2466.
- Bevacqua, Emanuele, Carl-Friedrich Schleussner, and Jakob Zscheischler**, “A Year above 1.5 °C Signals That Earth Is Most Probably within the 20-Year Period That Will Reach the Paris Agreement Limit,” *Nature Climate Change*, March 2025, *15* (3), 262–265.
- Bhargava, S. C.**, “Generalized Lotka-Volterra Equations and the Mechanism of Technological Substitution,” *Technological Forecasting and Social Change*, July 1989, *35* (4), 319–326.
- Bowen, Alex and Cameron Hepburn**, “Green Growth: An Assessment,” *Oxford Review of Economic Policy*, October 2014, *30* (3), 407–422.

- Bunn, Derek W. and Ahti A. Salo**, “Forecasting with Scenarios,” *European Journal of Operational Research*, August 1993, 68 (3), 291–303.
- Byers, Edward, Volker Krey, Elmar Kriegler, Keywan Riahi, Roberto Schaeffer, Jarmo Kikstra, Robin Lamboll, Zebedee Nicholls, Marit Sandstad, Chris Smith, van der Wijst, Kaj, Al -Khourdajie, Alaa, Franck Lecocq, Portugal-Pereira, Joana, Yamina Saheb, Anders Stromman, Harald Winkler, Cornelia Auer, Elina Brutschin, Matthew Gidden, Philip Hackstock, Mathijs Harm- sen, Daniel Huppmann, Peter Kolp, Claire Lepault, Jared Lewis, Giacomo Marangoni, Müller-Casseres, Eduardo, Ragnhild Skeie, Michaela Wern- ing, Katherine Calvin, Piers Forster, Celine Guivarch, Tomoko Hasegawa, Malte Meinshausen, Glen Peters, Joeri Rogelj, Bjorn Samset, Julia Stein- berger, Massimo Tavoni, and van Vuuren, Detlef**, “AR6 Scenarios Database,” <https://zenodo.org/record/5886911> November 2022.
- Bücker, Joris, R. Maria Del Rio-Chanona, Anton Pichler, Matthew C. Ives, and J. Doyne Farmer**, “Employment dynamics in a rapid decarbonization of the US power sector,” *Joule*, February 2025, 9 (2), 101803.
- Cannon, Alex J.**, “Twelve Months at 1.5 °C Signals Earlier than Expected Breach of Paris Agreement Threshold,” *Nature Climate Change*, March 2025, 15 (3), 266–269.
- Cherp, Aleh, Vadim Vinichenko, Jale Tosun, Joel A. Gordon, and Jessica Jewell**, “National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets,” *Nature Energy*, 2021, 6, 742–754.
- Clements, Michael P. and David F. Hendry, eds**, *A Companion to Economic Foreccacsting* Blackwell Companions to Sociology, Hoboken: Wiley, 2008.
- Comin, Diego, Bart Hobijn, and Emilie Rovito**, “Five Facts You Need to Know About Technology Diffusion,” *NBER Working Paper*, January 2006.

- Cowan, Robin**, “Tortoises and Hares: Choice Among Technologies of Unknown Merit,” *The Economic Journal*, 1991, 101 (407), 801–814.
- Creutzig, Felix, Peter Agoston, Jan Christoph Goldschmidt, Gunnar Luderer, Gregory F. Nemet, and Robert C. Pietzcker**, “The Underestimated Potential of Solar Energy to Mitigate Climate Change,” *Nature Energy*, August 2017, 2 (9).
- de Castro, Carlos, Margarita Mediavilla, Luis Javier Miguel, and Fernando Frechoso**, “Global Solar Electric Potential: A Review of Their Technical and Sustainable Limits,” *Renewable and Sustainable Energy Reviews*, December 2013, 28, 824–835.
- Dietz, Simon and Cameron Hepburn**, “Benefit–Cost Analysis of Non-Marginal Climate and Energy Projects,” *Energy Economics*, November 2013, 40, 61–71.
- Dixit, Robert K. and Robert S. Pindyck**, *Investment under Uncertainty*, Princeton University Press, December 1994.
- Ellis, Geraint, Nina Schneider, and Rolf Wüstenhagen**, “Dynamics of Social Acceptance of Renewable Energy: An Introduction to the Concept,” *Energy Policy*, October 2023, 181, 113706.
- Ember**, “Yearly Electricity Data,” <https://ember-climate.org/data-catalogue/yearly-electricity-data/> 2024.
- Farmer, J. Doyne**, *Making Sense of Chaos: A Better Economics for a Better World*, Dublin: Allen Lane, 2024.
- **and Arjun Makhijani**, “A US Nuclear Future? Not Wanted, Not Needed,” *Nature*, September 2010, 467 (7314), 391–393.
- **and François Lafond**, “How Predictable Is Technological Progress?,” *Research Policy*, April 2016, 45 (3), 647–665.

- , Cameron Hepburn, Matthew C. Ives, Thomas Hale, Thom Wetzer, Penny Mealy, Ryan Rafaty, Sugandha Srivastav, and Rupert Way, “Sensitive Intervention Points in the Post-Carbon Transition,” *Science*, April 2019, *364* (6436), 132–134.
- , – , Penny Mealy, and Alexander Teytelboym, “A Third Wave in the Economics of Climate Change,” *Climate Research in Gothenburg on*, 2015, *62*, 329–357.
- Feroli, Francesco, Koen Schoots, and Bob van der Zwaan, “Use and Limitations of Learning Curves for Energy Technology Policy: A Component-Learning Hypothesis,” *Energy Policy*, July 2009, *37* (7), 2525–2535.
- Funk, Jeffrey L. and Christopher L. Magee, “Rapid Improvements with No Commercial Production: How Do the Improvements Occur?,” *Research Policy*, April 2015, *44* (3), 777–788.
- Gambhir, Ajay, Richard Green, Michael Grubb, Philip Heptonstall, Charlie Wilson, and Robert Gross, “How Are Future Energy Technology Costs Estimated? Can We Do Better?,” *International Review of Environmental and Resource Economics*, December 2021, *15* (4), 271–318.
- Girod, Bastien, Arnim Wiek, Harald Mieg, and Mike Hulme, “The Evolution of the IPCC’s Emissions Scenarios,” *Environmental Science & Policy*, April 2009, *12* (2), 103–118.
- Gittins, John C., Richard Weber, and Kevin D. Glazebrook, *Multi-Armed Bandit Allocation Indices*, second ed., Hoboken, NJ: John Wiley & Sons, 2011.
- Grubb, Michael, Paul Drummond, Jean-Francois Mercure, Cameron Hepburn, Peter Barbrook-Johnson, Carlos Fjoao Joao Carlos Ferraz, Alex Clark, Laura Diaz Anadon, J. Doyne Farmer, Ben Hinder, Matthew C. Ives, Aled Jones, Gao Jun, Ulka Kelkar, Sergey Kolesnikov, Aileen Lam, Ritu Mathur,

**Roberto Pasqualino, Cristina Penasco, Hector Pollitt, Luma Ramos, Andrea Roventini, Pablo Salas, Simon Sharpe, Zhu Songli, Pim Vercoulen, Kamna Waghray, and Zhang Xiliang**, “The New Economics of Innovation and Transition: Evaluating Opportunities and Risks,” Technical Report, EEIST 2021.

**Grübler, Arnulf**, *The Rise and Fall of Infrastructures: Dynamics of Evolution and Technological Change in Transport*, Heidelberg: Physica-Verlag, 1990.

– , **Nebojša Nakićenović, and David G. Victor**, “Dynamics of Energy Technologies and Global Change,” *Energy Policy*, May 1999, 27 (5), 247–280.

**Hallegatte, Stéphane, Catrina Godinho, Jun Rentschler, Paolo Avner, Ira Irina Dorband, Camilla Knudsen, Jana Lemke, and Penny Mealy**, *Within Reach: Navigating the Political Economy of Decarbonization (Advance Edition)*, The World Bank, November 2023.

**Hsieh, I-Yun Lisa, Menghsuan Sam Pan, Yet-Ming Chiang, and William H. Green**, “Learning Only Buys You so Much: Practical Limits on Battery Price Reduction,” *Applied Energy*, April 2019, 239, 218–224.

**IEA**, “Electricity 2024 - Analysis and Forecast to 2026,” Technical Report May 2024.

– , “World Energy Investment 2024,” 2024.

**Ingersoll, Jonathan E.**, *Theory of Financial Decision Making* Rowman & Littlefield Studies in Financial Economics, third (print) ed., Savage, MD: Rowman & Littlefield, 1987.

**IPCC**, *Climate Change 2022, Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK and New York, NY, USA: IPCC, 2022.

- IPCC**, *Climate Change 2022, Impacts, Adaptation and Vulnerability: Working Group II Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 1 ed., Cambridge University Press, June 2023.
- IRENA**, “Renewable Power Generation Costs in 2021,” Technical Report, IRENA 2021.
- Jacquet-Lagrèze, Eric and Yannis Siskos**, “Preference Disaggregation: 20 Years of MCDA Experience,” *European Journal of Operational Research*, April 2001, *130* (2), 233–245.
- Kavlak, Goksin, James McNerney, and Jessika E. Trancik**, “Evaluating the Causes of Cost Reduction in Photovoltaic Modules,” *Energy Policy*, December 2018, *123*, 700–710.
- Kelleher, John D., Brian Mac Namee, and Aoife D’Arcy**, *Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies* The MIT Press Ser, second ed., Cambridge, Massachusetts: The MIT Press, 2020.
- Lafond, François, Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J. Doyne Farmer**, “How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts,” *Technological Forecasting and Social Change*, March 2018, *128*, 104–117.
- , **Diana Seave Greenwald, and J. Doyne Farmer**, “Can Stimulating Demand Drive Costs Down? World War II as a Natural Experiment,” *SSRN Electronic Journal*, June 2020.
- Lotka, Alfred J.**, *Elements of Mathematical Biology*, New York State: Dover Publications, 1956.
- Markard, Jochen**, “The next Phase of the Energy Transition and Its Implications for Research and Policy,” *Nature Energy*, August 2018, *3* (8), 628–633.

- McLaren, Duncan and Nils Markusson**, “The Co-Evolution of Technological Promises, Modelling, Policies and Climate Change Targets,” *Nature Climate Change*, May 2020, *10* (5), 392–397.
- McNerney, James, J. Doyne Farmer, Sidney Redner, and Jessika E. Trancik**, “Role of Design Complexity in Technology Improvement,” *Proceedings of the National Academy of Sciences*, May 2011, *108* (22), 9008–9013.
- Meade, Nigel and Towhidul Islam**, “Technological Forecasting – Model Selection, Model Stability, and Combining Models,” *Management Science*, 1998, *44* (8), 1115–1130.
- Mealy, Penny and Alexander Teytelboym**, “Economic Complexity and the Green Economy,” *Research Policy*, October 2022, *51* (8), 103948.
- , **Pete Barbrook-Johnson, Matthew C. Ives, Sugandha Srivastav, and Cameron Hepburn**, “Sensitive Intervention Points: A Strategic Approach to Climate Action,” *Oxford Review of Economic Policy*, November 2023, *39* (4), 694–710.
- Meng, Jing, Rupert Way, Elena Verdolini, and Laura Diaz Anadon**, “Comparing Expert Elicitation and Model-Based Probabilistic Technology Cost Forecasts for the Energy Transition,” *Proceedings of the National Academy of Sciences of the United States of America*, July 2021, *118* (27).
- Mercure, Jean-François, Pablo Salas, Pim Vercoulen, Gregor Semieniuk, Aileen Lam, Hector Pollitt, Phil B. Holden, Negar Vakilifard, Unnada Chewpreecha, Neil R. Edwards, and Jorge E. Vinales**, “Reframing Incentives for Climate Policy Action,” *Nature Energy*, December 2021, *6* (12), 1133–1143.
- Mulugetta, Yacob, Youba Sokona, Philipp A. Trotter, Samuel Fankhauser, Jessica Omukuti, Lucas Somavilla Croxatto, Bjarne Steffen, Meron Tesfamichael, Edo Abraham, Jean-Paul Adam, Lawrence Agbemabiese, Churchill Agutu, Mekalia Paulos Aklilu, Olakunle Alao, Bothwell Batidzirai, Getachew Bekele,**

- Anteneh G. Dagnachew, Ogunlade Davidson, Fatima Denton, E. Ogheneruona Diemuodeke, Florian Egli, Gebrekidan Gebresilassie Eshetu, Muluaem Gebreslassie, Mamadou Goundiam, Haruna Kachalla Gujba, Yohannes Hailu, Adam D. Hawkes, Stephanie Hirmer, Helen Hoka, Mark Howells, Abdulrasheed Isah, Daniel Kammen, Francis Kemausuor, Ismail Khennas, Wikus Kruger, Ifeoma Malo, Linus Mofor, Minette Nago, Destenie Nock, Chukwumerije Okereke, S. Nadia Ouedraogo, Benedict Probst, Maria Schmidt, Tobias S. Schmidt, Carlos Shenga, Mohamed Sokona, Jan Christoph Steckel, Sebastian Sterl, Bernard Tembo, Julia Tomei, Peter Twesigye, Jim Watson, Harald Winkler, and Abdulmutalib Yussuff, “Africa Needs Context-Relevant Evidence to Shape Its Clean Energy Future,” *Nature Energy*, October 2022, 7, 1015–1022.
- Nagy, Bela, J. Doayne Farmer, Quan M. Bui, and Jessika E. Trancik, “Statistical Basis for Predicting Technological Progress,” *PLoS ONE*, February 2013, 8 (2), 52669.
- Nature Energy, ed., “Uncertainty amid Political Changes,” *Nature Energy*, September 2024, 9 (9), 1043–1043.
- Neij, Lena, “Cost Development of Future Technologies for Power Generation—A Study Based on Experience Curves and Complementary Bottom-up Assessments,” *Energy Policy*, June 2008, 36 (6), 2200–2211.
- Nemet, Gregory F., “Beyond the Learning Curve: Factors Influencing Cost Reductions in Photovoltaics,” *Energy Policy*, November 2006, 34 (17), 3218–3232.
- , Jenna Greene, Finn Müller-Hansen, and Jan C. Minx, “Dataset on the Adoption of Historical Technologies Informs the Scale-up of Emerging Carbon Dioxide Removal Measures | Communications Earth & Environment,” *Communications Earth & Environment*, October 2023, 4 (1), 397.

– , **Laura Diaz Anadon, and Elena Verdolini**, “Quantifying the Effects of Expert Selection and Elicitation Design on Experts’ Confidence in Their Judgments About Future Energy Technologies,” *Risk Analysis*, 2017, *37* (2), 315–330.

**Nordhaus, William D.**, “An Optimal Transition Path for Controlling Greenhouse Gases,” *Science*, 1992, *258* (5086), 1315–1319.

– , “The Perils of the Learning Model for Modeling Endogenous Technological Change,” *Energy Journal*, January 2014, *35* (1), 1–13.

**Odenweller, Adrian, Falko Ueckerdt, Gregory F. Nemet, Miha Jensterle, and Gunnar Luderer**, “Probabilistic Feasibility Space of Scaling up Green Hydrogen Supply,” *Nature Energy*, September 2022, *7*, 854–865.

**Olivetti, Elsa A., Gerbrand Ceder, Gabrielle G. Gaustad, and Xinkai Fu**, “Lithium-Ion Battery Supply Chain Considerations: Analysis of Potential Bottlenecks in Critical Metals,” *Joule*, October 2017, *1* (2), 229–243.

**OWID**, “Energy Use per Person vs. GDP per Capita,” <https://ourworldindata.org/grapher/energy-use-per-person-vs-gdp-per-capita> 2023.

**Petropoulos, Fotios, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K. Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J. Bessa, Jakub Bijak, John E. Boylan, Jethro Browell, Claudio Carnevale, Jennifer L. Castle, Pasquale Cirillo, Michael P. Clements, Clara Cordeiro, Fernando Luiz Cyrino Oliveira, Shari De Baets, Alexander Dokumentov, Joanne Ellison, Piotr Fiszeder, Philip Hans Franses, David T. Frazier, Michael Gilliland, M. Sinan Gönül, Paul Goodwin, Luigi Grossi, Yael Grushka-Cockayne, Mariangela Guidolin, Massimo Guidolin, Ulrich Gunter, Xiaojia Guo, Renato Guseo, Nigel Harvey, David F. Hendry, Ross Hollyman, Tim Januschowski, Jooyoung Jeon, Victor Richmond R. Jose, Yanfei Kang, Anne B. Koehler, Stephan Kolassa, Nikolaos Kourentzes, Sonia Leva, Feng**

Li, Konstantia Litsiou, Spyros Makridakis, Gael M. Martin, Andrew B. Martinez, Sheik Meeran, Theodore Modis, Konstantinos Nikolopoulos, Dilek Önkal, Alessia Paccagnini, Anastasios Panagiotelis, Ioannis Panapakidis, Jose M. Pavía, Manuela Pedio, Diego J. Pedregal, Pierre Pinson, Patrícia Ramos, David E. Rapach, J. James Reade, Bahman Rostami-Tabar, Michał Rubaszek, Georgios Sermpinis, Han Lin Shang, Evangelos Spiliotis, Aris A. Syntetos, Priyanga Dilini Talagala, Thiyanga S. Talagala, Len Tashman, Dimitrios Thomakos, Thordis Thorarinsdottir, Ezio Todini, Juan Ramón Trapero Arenas, Xiaoqian Wang, Robert L. Winkler, Alisa Yusupova, and Florian Ziel, “Forecasting: Theory and Practice,” *International Journal of Forecasting*, July 2022, *38* (3), 705–871.

Pfeiffer, Alexander, Cameron Hepburn, Adrien Vogt-Schilb, and Ben Caldecott, “Committed Emissions from Existing and Planned Power Plants and Asset Stranding Required to Meet the Paris Agreement,” *Environmental Research Letters*, May 2018, *13* (5), 054019.

Pindyck, Robert S., “Climate Change Policy: What Do the Models Tell Us?,” *NBER Working Paper*, July 2013.

Plazas-Niño, Fernando Antonio, Nestor Raul Ortiz-Pimiento, and Erik Giovany Montes-Páez, “National Energy System Optimization Modelling for Decarbonization Pathways Analysis: A Systematic Literature Review,” *Renewable and Sustainable Energy Reviews*, July 2022, *162*, 112406.

Powell, Warren B., *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Hoboken, US: John Wiley & Sons, Incorporated, 2011.

Ramsey, Frank P., “A Mathematical Theory of Saving,” *The Economic Journal*, 1928, *38* (152), 543–559.

- Ritchie, Hannah**, “A Number of Countries Have Decoupled Economic Growth from Energy Use, Even If We Take Offshored Production into Account,” <https://ourworldindata.org/energy-gdp-decoupling> November 2021.
- , **Pablo Rosado, and Max Roser**, “CO<sub>2</sub> and Greenhouse Gas Emissions,” *Our World in Data*, 2023.
- Rodrik, Dani**, “Green Industrial Policy,” *Oxford Review of Economic Policy*, 2014, 30 (3), 469–491.
- Rowe, Gene and George Wright**, “The Delphi Technique as a Forecasting Tool: Issues and Analysis,” *International Journal of Forecasting*, October 1999, 15 (4), 353–375.
- Sabatier, Paul A.**, “Top-Down and Bottom-Up Approaches to Implementation Research: A Critical Analysis and Suggested Synthesis,” *Journal of Public Policy*, January 1986, 6 (1), 21–48.
- Salo, Ahti, Michalis Doumpos, Juuso Liesiö, and Constantin Zopounidis**, “Fifty Years of Portfolio Optimization,” *European Journal of Operational Research*, October 2024, 318 (1), 1–18.
- Sinha, Aditya, Aranya Venkatesh, Katherine Jordan, Cameron Wade, Hadi Eshraghi, Anderson R. De Queiroz, Paulina Jaramillo, and Jeremiah X. Johnson**, “Diverse decarbonization pathways under near cost-optimal futures,” *Nature Communications*, 2024, 15 (1), 8165.
- Sinsel, Simon R., Jochen Markard, and Volker H. Hoffmann**, “How Deployment Policies Affect Innovation in Complementary Technologies—Evidence from the German Energy Transition,” *Technological Forecasting and Social Change*, December 2020, 161.
- Smil, Vaclav**, “Examining Energy Transitions: A Dozen Insights Based on Performance,” *Energy Research & Social Science*, December 2016, 22, 194–197.

- Sognaes, Ida**, “What Can We Learn from Probabilistic Feasibility Assessments?,” *Joule*, November 2022, *6* (11), 2450–2452.
- Sørensen, Bent**, “Energy and Resources,” *Science*, July 1975, *189* (4199), 255–260.
- Stern, Nicholas H.**, *The Economics of Climate Change: The Stern Review*, Cambridge: Cambridge University Press, 2014.
- Trutnevyte, Evelina**, “Does Cost Optimization Approximate the Real-World Energy Transition?,” *Energy*, July 2016, *106*, 182–193.
- van de Ven, Dirk-Jan, Iñigo Capellan-Peréz, Iñaki Arto, Ignacio Cazarro, Carlos de Castro, Pralit Patel, and Mikel Gonzalez-Eguino**, “The Potential Land Requirements and Related Land Use Change Emissions of Solar Energy,” *Scientific Reports*, February 2021, *11* (1), 2907.
- Volterra, V.**, *Una Teoria Matematica Sulla Lotta per l’esistenza*, Zanichelli, 1927.
- Wagenvoort, Benjamin, Joel Dyer, François Lafond, and J. Doyne Farmer**, “The Universality of Technology Production,” 2025.
- Way, Rupert, François Lafond, Fabrizio Lillo, Valentyn Panchenko, and J. Doyne Farmer**, “Wright Meets Markowitz: How Standard Portfolio Theory Changes When Assets Are Technologies Following Experience Curves,” *Journal of Economic Dynamics and Control*, April 2019, *101*, 211–238.
- , **Matthew C. Ives, Penny Mealy, and J. Doyne Farmer**, “Empirically Grounded Technology Forecasts and the Energy Transition,” *Joule*, September 2022, *6*, 1–26.
- Wilson, Charlie**, “Up-Scaling, Formative Phases, and Learning in the Historical Diffusion of Energy Technologies,” *Energy Policy*, June 2012, *50*, 81–94.
- World Meteorological Organization**, “WMO Confirms 2024 as Warmest Year on Record at about 1.55°C above Pre-Industrial Level,” <https://wmo.int/news/media-centre/wmo->

confirms-2024-warmest-year-record-about-155degc-above-pre-industrial-level January 2025.

**World Steel Association**, “Fact Sheet: Energy Use in the Steel Industry,” April 2021.

**Wright, Theodore Paul**, “Factors Affecting the Cost of Airplanes,” *Journal of the Aeronautical Sciences*, February 1936, 3 (4).

**Xiao, Mengzhu, Tobias Junne, Jannik Haas, and Martin Klein**, “Plummeting Costs of Renewables - Are Energy Scenarios Lagging?,” *Energy Strategy Reviews*, May 2021, 35, 100636.

**Xu, Ruochong, Dan Tong, Steven J. Davis, Xinying Qin, Jing Cheng, Qinren Shi, Yang Liu, Cuihong Chen, Liu Yan, Xizhe Yan, Huaxuan Wang, Dongsheng Zheng, Kebin He, and Qiang Zhang**, “Plant-by-Plant Decarbonization Strategies for the Global Steel Industry,” *Nature Climate Change*, October 2023, 13 (10), 1067–1074.

# Chapter II

## Literature review

This thesis lies at the intersection of energy economics and stochastic modelling. While many research questions originate in the energy transition, the methods used to solve these questions are interesting and relevant in their own right. This Chapter outlines the academic context of the research questions and methods included in this thesis. It provides background to the energy transition and the role of technological change therein before diving into the different methodological backgrounds needed to answer wider questions in this field that are addressed in this thesis.

### II.1 Technological change in the energy transition

The transition to a low-carbon energy system is one of the largest technological change stories in history. As a result, a vast literature has emerged to study the drivers of this technological change and its implications on our society and the environment. The IPCC's Working Group III AR6 report in 2022 on climate change mitigation alone was developed by 199 Coordinating and Lead Authors (IPCC, 2022). We, therefore, limit our attention to the literature that is most relevant to this thesis, being the deployment and cost of clean energy technologies.

Clean technology deployment and cost have been central discussion topics since their inception. As early as 1909, newspaper reports detailed the cost of the newly invented solar panels, estimating a price equivalent to 660 USD/W in today's terms (Srivastav, 2023).<sup>1</sup> By the 1970s, comprehensive renewable energy scenarios emerged, such as those presented by Sørensen (1975) and Lovins (1977), which described 100% renewable energy pathways for Denmark and the United States. The oil crisis further spurred these discussions,

---

<sup>1</sup>Unfortunately, it would take many more decades for solar panel technology to take off, in part due to the kidnapping of its inventor (Srivastav, 2023).

highlighting energy costs and security concerns (Breyer et al., 2022). In 1992, the first IPCC report considered high-penetration scenarios, with projections indicating that up to 43% of total energy could be derived from nuclear and renewables by 2100 (Leggett et al., 1992).

### II.1.1 Deployment of clean energy technologies

Much of today’s debate on clean technology deployment revolves around two primary themes: the required penetration and growth rates (target setting) and the limitations thereof. Forecasting future deployment is challenging due to the numerous factors influencing technological adoption. Historically, some technologies’ growth potential has frequently been underestimated, like solar photovoltaics (Creutzig et al., 2017). Others, such as carbon capture and storage, have been subject to overly optimistic projections regarding their expansion (Mac Dowell et al., 2017).

A substantial body of research has examined the techno-economic constraints of individual clean energy technologies from a statistical perspective. Since these studies infer future deployment from historical data, they effectively serve as forecasts (Shmueli, 2010; Clements and Hendry, eds, 2008). This includes studies that propose hard or soft limits to deployment based on past growth data (Wilson, 2012; van Sluisveld et al., 2015; Wilson et al., 2013; Hansen et al., 2017; Höök et al., 2012; Madsen and Hansen, 2019; Napp et al., 2017; Odenweller et al., 2022; Cherp et al., 2021; Iyer et al., 2015; Odenweller and Ueckerdt, 2025; Kazlou et al., 2024; Reiner, 2016; Nemet et al., 2023; Link et al., 2025). However, as discussed in Chapter III, many of these findings are unreliable.

Other studies assess deployment constraints from physical and social perspectives. For example, Nonhebel (2005) assesses land-use constraints for solar- and bioenergy and Grant et al. (2022) investigate geological constraints to carbon capture and storage. Identifying deployment constraints ‘bottom-up’ is an inherently complex task given the evolving nature of technologies and societal norms—what may once have been a hard boundary is no longer because the underlying technology has changed. For example, land availability

was once seen as a considerable constraint to solar expansion. However, recent discussions have shifted toward coordination and environmental impacts. Similarly, battery markets are currently considered constrained by the availability of raw materials and their socio-economic implications. This may, however, change quickly if battery chemistry evolves (Olivetti et al., 2017; van de Ven et al., 2021; Grant et al., 2022; Lane et al., 2021; Dooley, 2013; Capellán-Pérez et al., 2017; de Castro et al., 2013; Hernandez et al., 2015; Valero et al., 2018; Turkovska et al., 2024; Diffendorfer et al., 2024). On top of that, it is essential to contextualize deployment bottlenecks within the current and historic energy system. For example, past projections incorrectly forecast a peak in global oil supply during the 1970s. Ultimately, this did not materialize such that fossil fuel supply chains continue to consume substantial land and material resources (Ritchie, 2022; Nijmens et al., 2023).

On top of individual technology targets and their limitations, it is also important to evaluate the interdependencies of multiple competing and complementary technologies (Sinsel et al., 2020). Ramping up multiple technologies simultaneously comes with its own set of challenges, including the engineering of novel energy systems (Breyer et al., 2022), building up new supply chains (Mealy and Teytelboym, 2022), or managing skill shortages and labour market bottlenecks (Bücker et al., 2025). The interaction of multiple S-curves through technology ecosystems constitutes a research field in its own right that far exceeds the scope of this thesis (Adner and Kapoor, 2016; Gallagher et al., 2012).

Lastly, corporate and government targets based on future energy scenarios navigate the complexities of this ecosystem despite limited model applicability, as well as risks of inconsistency and greenwashing amid large uncertainties (Pielke and Ritchie, 2021; Pindyck, 2013; Blondeel et al., 2024). Going forward, steering the increasing complexities of clean technology deployment and the energy transition as a whole will require more targeted model development, as well as closer interaction with decision-makers to fit their needs (Barbrook-Johnson et al., 2024; Strachan et al., 2016).

## II.1.2 Cost of clean energy technologies

A parallel line of research investigates the historical and future costs of clean energy technologies. These analyses are also challenging due to the dynamic nature of technological costs, which are heterogeneous across components, vary by location, and depend on the broader innovation ecosystems (Gallagher et al., 2012).

Historical cost analyses are, by definition, limited to mature technologies such as solar photovoltaics and wind energy. Due to their importance in regional and global energy systems, there is a large literature analysing the historical evolution of these technologies. Table II.1 provides a non-comprehensive overview of recent studies on the topic. Due to the importance of regional costs in Chapter IV, we focus on studies with regional disaggregation of costs. For this thesis, Elshurafa et al. (2018) is worth highlighting, who perform an ordinary and seemingly unrelated regression analysis on drivers of national balance-of-system cost for solar in 20 countries that includes a local learning curve. While our results in Chapter IV contradict some of Elshurafa et al. (2018), this work was an important inspiration. On top of regional cost studies, there are important global cost analyses, particularly Nemet (2006); Kavlak et al. (2018), which investigate historical drivers of investment cost improvement in photovoltaic modules, Elia et al. (2020), which examine historical drivers of investment cost improvements in wind turbines and Steffen (2020), which assess the impact of financial conditions on renewable costs.

Many of these studies make implicit forecasts about future renewable costs. In addition, there are a number of studies with explicit forecasts for future costs. Next to expert elicitations (see Meng et al. (2021) and citations therein) and mechanistic models (Junginger et al., 2004; Chang et al., 2022), statistical models are commonly used to make forecasts (Farmer et al., 2015; Lafond et al., 2018; Grafström and Poudineh, 2021; Hayward and Graham, 2013; Wiesenthal et al., 2012; Rubin et al., 2015; Neij and Nemet, 2022; Neij, 2008; Glenk et al., 2021; Söderholm and Sundqvist, 2007; Steffen et al., 2020; Weiss et al., 2019; Ek and Söderholm, 2010; Samadi, 2018; Schmidt et al., 2017). These are also

<b>Study</b>	<b>Technology</b>	<b>Coverage</b>	<b>Description</b>
<a href="#">Elshurafa et al. (2018)</a>	Solar PV	20 different countries	Calculation of learning rates for national solar PV balance-of-system cost
<a href="#">Vartiainen et al. (2020)</a>	Solar PV	6 European locations	Effect of solar irradiation and WACC on LCOE
<a href="#">Kothari et al. (2023)</a>	Solar PV	13 different countries	Multi-factor (30) model for national solar PV balance-of-system cost
<a href="#">Barbose et al. (2015); Trancik et al. (2020)</a>	Solar PV	US	Investigation of historic drivers of LCOE
<a href="#">Strupeit and Neij (2017)</a>	Distributed solar PV	Germany	Investigation of historic drivers of investment costs
<a href="#">Feldman et al. (2021)</a>	Solar PV	US	Investigation of historic drivers of investment cost
<a href="#">Wiser et al. (2020)</a>	Solar PV	US	Calculation of learning curves for (O&M)
<a href="#">Apostoleris et al. (2018)</a>	Solar PV	Middle east	Investigation of historic drivers of low solar PV LCOE in the Middle East
<a href="#">Rai et al. (2020); Nemet et al. (2020)</a>	Solar PV	US	Investigation of spillovers and cost reduction in soft costs
<a href="#">Klemun et al. (2023)</a>	Solar PV	Global, US, Germany, Japan	Investigation of cost reduction in hard and soft investment cost
<a href="#">Bolinger et al. (2022)</a>	Onshore wind	US	Investigation of historic drivers of LCOE
<a href="#">Wiser et al. (2019)</a>	Onshore wind	US	Calculation of learning curves for operation & maintenance (O&M)
<a href="#">Riva et al. (2018)</a>	Onshore wind	Denmark, Germany, Ireland, Norway, Sweden, EU, US	Calculation of operation & maintenance cost (O&M)
<a href="#">Hayashi et al. (2018)</a>	Wind	China	Investigation of the impact of experience on generation performance, turbine size, unit cost
<a href="#">Steffen et al. (2020)</a>	Solar PV and onshore wind	Germany	Calculation learning curves for operation & maintenance (O&M) cost in Germany
<a href="#">Steffen (2018)</a>	Solar PV and onshore wind	Germany	Calculating learning curves for cost of capital
<a href="#">Raupach-Sumiya et al. (2015)</a>	Solar PV and onshore wind	Germany, Japan	Comparison of renewable investment and operating costs between Japan and Germany

Table II.1: **Examples of regional LCOE and bottom-up cost studies (not comprehensive).**

the focus here, as discussed in the introduction. Statistical models can also be used to investigate the limits to future cost reductions, also referred to as floor costs (Way et al., 2022; Berckmans et al., 2017; Hsieh et al., 2019), and make inferences for emerging technologies based on analogous technologies (Sievert et al., 2024).

### II.1.3 Integrated Cost and Deployment models

Integrating technology cost and deployment dynamics is a challenging task due to the complex causal interactions between the two and limited historical data (Lafond et al., 2020). As a result, a vast amount of literature has emerged modelling these interactions in the context of the energy transition. The literature applies different names to describe model outcomes, such as ‘projections’, ‘scenarios’, or ‘simulations’, to describe the objective and confidence in the results (Bunn and Salo, 1993). Here, we provide a brief and non-comprehensive overview of modelling techniques and objectives and refer interested readers to the reviews cited below.

Popular techno-economic models integrating deployment and cost evolution include optimization- and equilibrium-based bottom-up energy system models, including energy-system modules in integrated assessment models (IAMs), agent-based models (ABMs) focusing on the interaction of heterogeneous decision-makers, or econometric models consolidating different empirical results. Reviews of this extensive literature can be found in (Hoffman and Wood, 1976; Jebaraj and Iniyar, 2006; Bhattacharyya and Timilsina, 2010; Lopion et al., 2018; Prina et al., 2020; Pfenninger et al., 2014; Krey, 2014; Krey et al., 2019; van Beek et al., 2020; Mercure et al., 2018).

These technology-focused models are frequently combined with macroeconomic models (Lehtila and Loulou, 2005), following the pioneering DICE model by Nordhaus (1992). This integration enables assessing the wider economic implications of technological change. Technology costs and deployment can have large effects on value chains (Stern and Xie, 2023) and play a vital role in industrial policy and trade strategies, shaping international partnerships and growth trajectories (Rodrik, 2014; Mealy and Teytelboym, 2022; Rosenow

and Mealy, 2024; Mealy and Hepburn, 2020; Bowen and Hepburn, 2014; Smulders et al., 2014; Altenburg and Rodrik, 2017; Fouquet, ed, 2019; Andres et al., 2023; Semieniuk et al., 2022). Integrated models have already influenced a wide range of international policies on energy transition, including UN climate negotiations and IPCC reports (IPCC and WMO, eds, 1992; WMO and IPCC, eds, 1990b,a; McLaren and Markusson, 2020; IPCC, ed, 2007; Byers et al., 2022; IPCC, 2023). Nevertheless, integrating technological change into macroeconomic models comes with its own challenges. Most prominently, equilibrium conditions inherently associate deviations from the status quo with higher costs, often neglecting endogenous improvements in clean technologies (Farmer et al., 2015). One way to resolve these challenges is rigorous statistical analysis of underlying technological change models.

## II.2 Empirical analysis of technological change

Researchers have long been interested in understanding and forecasting technological change (Farmer, 2024; Lafond, 2025). As a result, a broad body of literature has emerged, focusing on assessing technological transitions from a statistical perspective. These models provide the foundation for much contemporary research aimed at predicting the energy transition. In this section, we focus on two primary empirical dynamics: S-curves for technology deployment and experience curves for technology costs.

Beyond S-curves and experience curves, other data-based methods exist for studying innovation dynamics but are outside of the scope of this thesis. For example, experience curves can be studied through the lens of production functions to examine the role of experience and technological improvements in productivity gains (Elia et al., 2020; Irwin and Klenow, 1994; Gruber, 1998). Another method frequently used to investigate key technological innovations and knowledge spillovers is patent analysis (Nemet and Johnson, 2012; Nemet, 2012, 2009; Hötte and Jee, 2022; Pichler et al., 2020). These can also be connected to experience curves (Triulzi et al., 2020; Benson and Magee, 2015), providing a natural extension to the work presented in this thesis.

## II.2.1 S-curve models of technology deployment

S-curves describe the adoption of new technologies over time. Historically, technological adoption follows a consistent pattern: it starts from a very small base and remains small for an extended period. However, it grows approximately exponentially from the start and accelerates until it reaches significant penetration before eventually plateauing at its maximum market size. The reverse pattern applies to technologies that are phased out or replaced by newer alternatives.

S-curve diffusion has been explored across a wide range of technologies, including transportation, manufacturing, and even social institutions such as monasteries in the Middle Ages (Grübler, 1990; Grübler et al., 1999; Grübler and Nakićenović, 1996; Grübler, 1991; Comin and Hobijn, 2010; Comin et al., 2006; Grübler et al., 2016). While the scale and speed of adoption vary across technologies, the fundamental shape of the diffusion remains similar.

A key feature of S-curves is their highly non-linear nature. The rate of change (i.e., the derivative of the S-curve) begins near zero, reaches its maximum at the inflection point, and then decreases again as the S-curve approaches its asymptote. This pattern highlights different phases of technological diffusion. The early S-curve stages are characterized by a formative phase (Wilson, 2012; Bento and Wilson, 2016), shaped by research efforts and initial niche markets. Later, this is followed by industrialization and mass markets before an eventual phase-out (Grübler and Nakićenović, 1991).

Various mathematical models describe S-curves, including the Gompertz model (Cherp et al., 2021) and the Richards model (Marinakos, 2012). However, the most widely used model is the logistic growth model described in the introduction (Grübler, 1990; Fisher and Pry, 1971), which originates from biological growth models, such as bacterial population dynamics. This model is symmetric around the inflection point, where growth initially appears exponential but eventually levels off to a constant asymptote.

Next to the question of model choice, applying S-curves to real data also necessitates

explicit or implicit treatment of statistical noise. Real-world processes generally do not follow an S-curve exactly, but are subject to additional constraints and fluctuations that cannot be captured easily without overfitting the data. In many applications, we cannot say with statistical significance which model best describes an S-curve due to this noise. Furthermore, during early stages of a technology’s deployment, we cannot even differentiate with statistical significance an S-curve model from an exponential growth model (Wagenvoort et al., 2025; Lafond, 2025).<sup>2</sup> For practical applications, the model choice will depend on both statistical description of the data, the level of noise present, as well as the question the model is meant to answer. For example, if the question is to identify the average historical growth rate, an exponential model may perform better than an S-curve model due to less risk of overfitting. However, if the task focusses on long-term scenario generation, an exponential model is limited since a technology cannot diffuse indefinitely.

The presence of noise also makes forecasting with S-curves challenging (Wagenvoort et al., 2025). Several attempts at prediction have been made but have often failed, as seen in the case of mobile phone adoption (Boretos, 2008). To date, no general unbiased forecasting model has been published (Petropoulos et al., 2022).

## II.2.2 Experience curve models of technology cost

Experience curves are used to describe the cost evolution of technologies over time. Unlike S-curves, experience curves have been successfully applied to generate robust forecasts, particularly in the energy sector. Similarly to S-curves, no technology follows an experience curve exactly, but its approximation can be useful.

The concept of experience curves was first introduced by Wright (1936), who described a decline in airplane manufacturing costs. Experience curves are, therefore, commonly referred to as Wright’s law. Wright observed that with each doubling of cumulative

---

<sup>2</sup>Note that an S-curve model will always have a higher  $R^2$ -value than an exponential model since it has at least one additional parameter. This, however, does not mean that it is a statistically significant better description of reality since overfitting may be present.

production, airplane costs declined by a fixed proportion, approximately following a power-law relationship. Another widely recognized cost curve is the learning curve, or Moore’s Law, named after Intel co-founder Gordon Moore (Moore, 1975). This model describes an exponential decline in a technology’s costs over time (Farmer and Lafond, 2016; Funk and Magee, 2015).

While these models effectively describe cost evolution, the underlying mechanisms are more nuanced and often heterogeneous (Grafström and Poudineh, 2021). Various factors influence cost reductions, including supply chain efficiency, economies of scale, manufacturing improvements, and research-driven innovation (Wilson, 2012; Kavlak et al., 2018). Other studies emphasize the role of design complexity and random innovation in cost declines (McNerney et al., 2011b), while some focus on reductions at the component level of a technology (Ferioli et al., 2009). As a result, more complex approaches have emerged from Wright’s and Moore’s law. Two-factor models, for example, integrate both Moore’s and Wright’s Laws or incorporate cost reductions driven by research and development expenditures (Nagy et al., 2013).

One of the most significant contributions of cost curve models is their ability to forecast future technological costs. Studies such as (Alberth, 2008; Nagy et al., 2013; Farmer and Lafond, 2016; Lafond et al., 2018) evaluate the predictive accuracy of different cost curves across a wide range of technologies. There are multiple findings from these studies worth discussing:

Firstly, the predictive performance of Wright’s and Moore’s is very similar, with neither model having a statistically significant advantage over the other (Nagy et al., 2013; Lafond et al., 2018). One reason for this is if a technology’s cumulative production grows approximately exponentially—a common occurrence in the early stages of an S-curve or in scenarios with elastic demand—both laws yield similar results (Ferioli and van der Zwaan, 2009; Nordhaus, 2014; Sahal, 1979). Nevertheless, choosing either Wright’s or Moore’s law can lead to drastically different policy implications. This limitation of either method means that it can be beneficial for modellers to investigate the robustness of their results

with respect to the model choice (Way et al., 2022).

Secondly, Nagy et al. (2013) find that more complex models, such as two-factor models, perform worse than their one-factor equivalent. This suggests that multi-factor models are subject to overfitting, leading to spurious regression results and poor out-of-sample performance (Clements and Hendry, eds, 2008). Lafond et al. (2020) examines this phenomenon in the context of World War II, where super-exponential production growth and inelastic demand revealed that Wright’s and Moore’s law contributed roughly equally to observed cost reductions. These findings imply that reality lies in between Moore’s and Wright’s law but a single factor model is still preferred for forecasting purposes due to lower risk of overfitting. Using a two-factor model in policy applications may nevertheless be justified but needs to be treated with care due to the likely overfitting.

Thirdly, no model so far has identified a specific technology or time period where cost curves deviated from Moore’s and Wright’s law with statistical significance. Moreover, previous models in the energy sector that assumed a floor cost for solar photovoltaics have consistently over-estimated future costs (Way et al., 2022). This implies that models assuming continuous cost evolution have outperformed alternatives with structural breaks or regime changes. It also suggests that continuous models are more applicable to forecast novel clean energy technology costs. Nevertheless, this does not mean that continuity will necessarily hold across technologies and time. As we will see in Chapter IV, we expect a regime change for onshore wind costs due to its underlying components. This may also occur in other technologies and for other reasons and is a key limitation for quantitative forecasting methods more widely.

Lastly, Farmer and Lafond (2016) and Lafond et al. (2018) were the first studies to explicitly model the stochastic variations around cost curves. They show that cost curves can be used for probabilistic forecasting and ensuring the robustness of predictions (Gluzberg and Katz, 2025), as well as forecasting future uncertainty in technology costs.

Building on Lafond et al. (2018), Way et al. (2022) constructs future energy system scenario costs by applying stochastic experience curves to clean energy technologies and

mean-reverting stochastic to forecast future costs of fossil fuel technologies (McNerney et al., 2011a). Their findings indicate that a rapid energy transition is more cost-effective than a slow transition or maintaining the status quo, even without accounting for climate change externalities. This demonstrates the importance of getting technological learning right in the predictions of energy systems models.

### II.2.3 Investment allocation models

Technology costs play an important role in determining their diffusion (IPCC, 2022). Likewise, technology cost forecasts—such as those based on Wright’s law—are useful in guiding resource allocation for developing future clean technologies, even if they are not yet cost-effective. Investors seek to identify profitable technology investment strategies, while operational leaders are frequently required to make supply chain commitments. Given the inherent uncertainty in future technology costs, resource allocation strategies have been a long-standing research area. In this section, we focus on key approaches, particularly discounted cash flow (DCF) analysis and associated optimization methods. We pay particular attention to non-linear and stochastic optimization, including novel machine-learning-based methods.

The predominant approach for calculating cost-optimal technology investments in the energy transition is DCF analysis. DCFs are widely used in both engineering-focused bottom-up models, such as capacity expansion models, and environmental economics models, such as various IAM modules (Trutnevyte, 2016; Conejo and Baringo, 2018; Wagner, 2014). In these models, investments are allocated to the portfolio of technologies that minimize the discounted sum of costs over a given time horizon or, equivalently, maximize the discounted sum of utility (Ingersoll, 1987). DCF models are particularly useful in identifying normative decarbonization pathways, guiding decision-making toward low-cost options (Trutnevyte, 2016). DCF models are also applied to cost-benefit analysis (CBA), where the costs of an environmental measure are weighed against its future societal benefits. This method is founded on economic equilibrium theory and marginal deviations

from equilibrium (Ramsey, 1928; Harberger, 1971; Arrow et al., 2013; Starrett, 1988).

Before discussing modelling techniques that optimise DCFs, it is worth highlighting key critiques and limitations of this approach.

Firstly, it is worth noting that CBA has faced increasing theoretical scrutiny. Key assumptions, such as the marginal impact of investments and their distributional effects, have been questioned in the context of the energy transition (Dietz and Hepburn, 2013; Ackerman and Heinzerling, 2002; Farrow, 2012; Grubb et al., 2021). As a result, alternative approaches are emerging, including risk-opportunity analysis (Grubb et al., 2021) and the study of sensitive intervention points (Farmer et al., 2019; Mealy et al., 2023; Winkelmann et al., 2022; Tàbara et al., 2018). These methods highlight that targeted political actions can have disproportionately large effects on both costs and environmental benefits.

Secondly, real-life decisions do not consistently follow theoretically optimal conditions (Trutnevyte, 2016). An emerging body of literature thus focuses on non-optimal decarbonization pathways that more accurately reflect real-world decision-making. These include agent-based models (ABMs) and econometric models, as well as sub-optimal models that explore the solution space around theoretically optimal decarbonization pathways (Farmer et al., 2015; Anwar et al., 2022; Grubb et al., 2021; Barbrook-Johnson et al., 2024; Jonson et al., 2020; Anwar et al., 2022; Yue et al., 2018; Sinha et al., 2024). Nevertheless, understanding cost-optimal pathways remains valuable for designing policies and investment strategies that drive clean technology adoption up and costs down.

Due to its wide use in research and practice, a vast number of optimization models have been developed to identify DCF-optimal investment allocation strategies. Most models share that they simulate future cash flows to identify cost-optimal outcomes (Yue et al., 2018; Herbst et al., 2012; Powell, 2011). Early models were mostly static and linear due to their lower computational requirements and simpler interpretation. More recently, linear models have been replaced by context optimization methods to incorporate non-linear effects, such as endogenous learning (Heuberger et al., 2017; Hayward and Graham, 2013; Louwen et al., 2019; Wagner, 2014). This required the introduction of new techniques,

such as mixed-integer optimization and partial linearization of the experience curve.

A second major direction in optimization modelling involves explicitly incorporating the uncertainty of future innovation (Webster et al., 2017; Labriet et al., 2015; Messner et al., 1996; Milford et al., 2022; Lopion et al., 2019; Vespucci et al., 2016; Hach and Spinler, 2018; Bukenberger and Palmintier, 2018). Methods such as multi-stage stochastic optimization and Monte Carlo analysis are commonly used, with a primary focus on parameter uncertainty. Other types of uncertainty, such as external shocks, can also be incorporated, though computational efficiency remains a key limitation of this approach. Including uncertainty in these models also allows for alternative investment objectives, particularly risk aversion. Investors seek not only to minimize costs but also to reduce the uncertainty associated with future costs (Markowitz, 1952, 2014; Way et al., 2019).

A third approach worth mentioning is myopic optimization (Keppo and Strubegger, 2010; Levin et al., 1983). Unlike perfect foresight models, which assume complete knowledge of future cash flows across the entire optimization horizon, myopic optimization considers decision-making over discrete time steps, often using a moving-window approach. While the resulting model is no longer strictly optimal, it aims to better capture real-world decision-making behaviour. However, combining this approach with other optimization objectives can also introduce undesired time inconsistencies (Rudloff et al., 2014).

For this thesis, particularly Chapter V, it is central to note that in practice, many of these models are applied dynamically. Instead of implementing a model once and acting on it, dynamic modelling involves the repeated application to refine the resource allocation throughout the energy transition (Samuelson, 1969; Webster et al., 2017; Powell, 2011; Cowan, 1991). The main challenge in dynamic portfolio optimization is the *curse of dimensionality*. The dynamic solution space, i.e., the number of possible dynamic strategies, is much larger than that of static strategies (Powell, 2011). Popular approaches to address this challenge include backward integration and decision tree approaches, although they are computationally intensive and can be difficult to interpret.

One dynamic portfolio optimization strategy that has made significant progress and

has been particularly successful in practice is reinforcement learning (Sutton and Barto, 2014). This approach has been widely applied in artificial intelligence and clinical research, allowing algorithms to simultaneously explore and optimize solutions. A central framework within reinforcement learning is the Multi-Armed Bandit (MAB) problem (Gittins, 1979; Gittins et al., 2011). Imagine a scenario where a gambler chooses between different slot machines (one-armed bandits) to maximise their rewards. Without knowing the underlying distributions, which machine should the gambler pick? The MAB framework models decision-making in situations where the reward distribution of different portfolio choices is initially unknown. The problem has been extensively studied and applied in recent decades (Gittins et al., 2011; Scott, 2010; Bank and Föllmer, 2003; Chakravorty and Mahajan, 2014).

There are two primary schools of thought for solving the MAB problem: index strategies and non-parametric strategies. Non-parametric strategies, such as Thompson sampling,  $\epsilon$ -greedy strategies, or upper confidence bounds, have been very successful in practical applications due to their reliability and ease of implementation. Even with little knowledge about the underlying uncertainty-generating processes, they can perform well (although not strictly optimal) (Agrawal and Goyal, 2012). Index strategies, such as the Gittins index and the Whittle index, provide asymptotically optimal solutions (Gittins, 1979; Gittins et al., 2011; Whittle, 1988). Their disadvantage in practical applications is often that they require exact knowledge of the underlying noise-generating processes. In addition, implementing these strategies can be quite involved.

In particular, implementing dynamic index strategies for portfolio optimization is closely connected to *real options theory* (Bank and Föllmer, 2003). Here, individual investment decisions are values as financial options would, for example, through a Black-Scholes model (Peskir and Shiryaev, 2006). Real options have previously been used to assess individual investment decisions in the energy transition (Anwar et al., 2022; Basei et al., 2024; Kumbaroğlu et al., 2008; Bruno et al., 2016; Fernandes et al., 2011). In Chapter V, we will see how similar methods also enable the study of long-term portfolio strategies

and the sensitive intervention points therein.

The arguably most important result of this literature is that, for dynamic decisions under uncertainty, higher volatility increases value (Dixit and Pindyck, 1994). This also applies to policy decisions, even if these are subject to additional constraints and path-dependencies. Risk-seeking behaviour can yield significant benefits in shaping policy for the energy transition (Rodrik, 2014).

## II.2.4 Validation of technological change models

Model validation is important across all three models examined in this thesis to improve the reliability of modelled results. Without rigorous validation, the credibility of the model outputs frequently remains unknown. However, as will be demonstrated in subsequent chapters, the process of validating a model can be as complex as than designing the model itself. Consequently, a substantial body of literature is dedicated to the topic of model validation. Broadly, validation can be categorized into empirical and theoretical approaches. Empirical validation assesses whether model results align with real-world data, whereas theoretical validation evaluates whether the results are self-consistent with the underlying assumptions used to construct the model (Petropoulos et al., 2022; Powell, 2011).

Before undertaking a review of empirical and theoretical validation methodologies, it is useful to clarify the primary objectives of model validation. Model validation can serve multiple purposes, though not all are equally relevant in the context of this thesis. The two most prominent purposes are model comparisons and reliability assessments (Fildes and Makridakis, 1995; Stephenson and Jolliffe, eds, 2012). Model comparisons involve evaluating the accuracy of different models against one another. Reliability assessments focus on ensuring the internal consistency of a single model. For example, we can compare the accuracy of two competing forecasting models. Alternatively, a forecasting model that predicts that future empirical values fall within a given range can be validated for its reliability (Farmer and Lafond, 2016). This thesis primarily focuses on model reliability.

Empirical validation techniques vary depending on the specific model and data used. In the context of S-curves, experience curves, and integrated cost and deployment models, several key aspects of empirical validation warrant attention.

First, validation can be conducted using *in-sample* and *out-of-sample* testing. In-sample validation involves calibrating and testing a model using the same dataset, while out-of-sample validation employs separate datasets for calibration and testing (Tashman, 2000). In-sample validation is commonly applied in social sciences and econometrics due to its simpler implementation, often using  $R^2$ -values or statistical significance tests for model parameters (Heij, ed, 2004). However, this approach can be misleading, as it typically necessitates strong assumptions about the noise structure of the data and poses a risk of overfitting (Tashman, 2000; Pant and Starbuck, 1990). Out-of-sample validation, while less dependent on these assumptions, still requires careful implementation to mitigate sample biases and information leakage (Wang and Ruf, 2022). Additionally, out-of-sample validation demands significantly more data, as adequate sample sizes are needed for both training and testing purposes.

Both in-sample and out-of-sample validation often utilise visual analysis of data. While this is certainly helpful, quantitative analysis of model validation is often preferred, a practice known as *scoring rules* (Gneiting and Raftery, 2007; Murphy and Winkler, 1987). A scoring rule quantifies a model's accuracy or reliability for an empirical sample by projecting the model performance to a single number. Among the most commonly employed scoring rules in social science literature is the  $R^2$ -value. However, many alternative scoring rules exist, such as mean and median absolute errors. The choice of scoring rule is contingent on the nature of the forecast and the assumed error model (Stephenson and Jolliffe, eds, 2012). For example, the mean absolute error is preferred over the  $R^2$ -value in the presence of heavy-tailed error distributions since it is more robust to large errors (Chai and Draxler, 2014). When forecasts generate intervals or distributions of future values, more complex scoring rules are required that capture these outputs. One such technique is the assessment of forecast intervals using unconditional coverage (Gneiting and Raftery, 2007). This

can be further generalized to distributional forecasts through the probability integral transformation, which evaluates both accuracy and confidence simultaneously (Gneiting et al., 2007; Gneiting and Katzfuss, 2014).

A critical challenge in all empirical exercises is assessing their statistical significance to identify whether an empirical result was driven by noise. Based on the underlying model, statistical significance can be evaluated either analytically or heuristically using surrogate data (Petropoulos et al., 2022). Significance testing requires additional data and can be constrained by data correlation. If data is correlated, its noise relative to a model is not independent, lowering the effective degrees of freedom (Bretherton et al., 1999). This data limitation and statistical significance challenge is particularly restrictive when noise and uncertainty themselves are endogenous. *Endogenous uncertainty* refers to cases where model results impact statistical noise (Salo et al., 2024). Chapter V explores such a situation in detail. Under Wright’s law, a technology’s cost changes with its cumulative deployment, subject to stochastic noise. This means that both costs and the associated uncertainty can be seen as endogenous with respect to the technology’s deployment. With endogenous uncertainty, deployment model results cannot be tested empirically using ex-post data. Instead, it requires us to actually use the model to test it.

In some cases, data constraints can be relieved by applying theoretical validation methods. Theoretical validation, unlike empirical validation, does not rely on empirical data. Theoretical models can serve as priors in selecting appropriate models, contingent on the available data. Analytical analyses, for instance, can indicate which model choices are expected to minimize forecast errors (Smart and Farmer, 2024; Pesaran et al., 2024; Grunfeld and Griliches, 1960; Hendry and Hubrich, 2011). As more data becomes available, these analytical priors may evolve to include new insights. Similarly, heuristic analyses offer insights into model selection by simulating different alternatives and establishing priors based on simulated performance (Scott, 2010; Breiman, 1996; Agrawal and Goyal, 2012). Heuristics usually cannot replace analytical assessments but may be sufficient in practical applications.

Model selection inherently depends on a set of assumptions that must hold for a given model to be appropriate. Forecasting models, for example, often assume stationarity, while optimization models assume specific error structures. Theoretical validation exercises facilitate a comprehensive evaluation of these assumptions, highlighting the conditions under which deviations might lead to significantly different outcomes ([Markowitz, 2014](#); [Way et al., 2019](#); [Cowan, 1991](#)). Moreover, theoretical arguments can be employed to identify equivalent models. For instance, [Nordhaus \(2014\)](#) demonstrates that in the presence of demand elasticity, Moore’s Law and Wright’s Law models are equivalent, implying that these models cannot be distinguished empirically under elastic demand conditions.

In summary, model validation is a critical step in improving the reliability of technological change models. Empirical validation methods, such as in-sample and out-of-sample testing, provide data-driven assessments of model accuracy, while theoretical validation techniques offer insights where empirical validation is constrained by data limitations.

## References

- Ackerman, Frank and Lisa Heinzerling**, “Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection,” *University of Pennsylvania Law Review*, 2002, 150 (5), 1553–1584.
- Adner, Ron and Rahul Kapoor**, “Innovation Ecosystems and the Pace of Substitution: Re-examining Technology S-curves,” *Strategic Management Journal*, 2016, 37 (4), 625–648.
- Agrawal, Shipra and Navin Goyal**, “Analysis of Thompson Sampling for the Multi-armed Bandit Problem,” *JMLR: Workshop and Conference Proceedings*, 2012, 23, 39.1–39.26.
- Alberth, Stephan**, “Forecasting Technology Costs via the Experience Curve — Myth or Magic?,” *Technological Forecasting and Social Change*, September 2008, 75 (7), 952–983.
- Altenburg, Tilman and Dani Rodrik**, *Green Industrial Policy: Accelerating Structural Change towards Wealthy Green Economies*, Geneva, Bonn: UN Environment; German Development Institute / Deutsches Institut für Entwicklungspolitik (DIE), 2017.
- Andres, Pia, Penny Mealy, Nils Handler, and Samuel Fankhauser**, “Stranded Nations? Transition Risks and Opportunities towards a Clean Economy,” *Environmental Research Letters*, April 2023, 18 (4), 045004.
- Anwar, Muhammad Bashar, Gord Stephen, Sourabh Dalvi, Bethany Frew, Sean Ericson, Maxwell Brown, and Mark O’Malley**, “Modeling Investment Decisions from Heterogeneous Firms under Imperfect Information and Risk in Wholesale Electricity Markets,” *Applied Energy*, January 2022, 306, 117908.
- Apostoleris, Harry, Sgouris Sgouridis, Marco Stefancich, and Matteo Chiesa**, “Evaluating the Factors That Led to Low-Priced Solar Electricity Projects in the Middle East,” *Nature Energy*, December 2018, 3, 1109–1114.

**Arrow, Kenneth, Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, Robert S. Pindyck, William Pizer, Paul Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman**, “Determining Benefits and Costs for Future Generations,” *Science*, July 2013, *341* (6144), 349–350.

**Bank, Peter and Hans Föllmer**, “American Options, Multi-Armed Bandits, and Optimal Consumption Plans: A Unifying View,” *SFB 373 Discussion Paper*, 2003, *2003* (46), 1–42.

**Barbose, Galen, Naïm Richard Darghouth, Samantha Weaver, David Feldman, Robert Margolis, and Ryan Wiser**, “Tracking US Photovoltaic System Prices 1998–2012: A Rapidly Changing Market,” *Progress in Photovoltaics: Research and Applications*, 2015, *23* (6), 692–704.

**Barbrook-Johnson, Pete, Jean-François Mercure, Simon Sharpe, Cristina Peñasco, Cameron Hepburn, Laura Diaz Anadon, J. Doyne Farmer, and Timothy M. Lenton**, “Economic Modelling Fit for the Demands of Energy Decision Makers,” *Nature Energy*, March 2024, *9* (3), 229–231.

**Basei, Matteo, Giorgio Ferrari, and Neofytos Rodosthenous**, “Uncertainty over Uncertainty in Environmental Policy Adoption: Bayesian Learning of Unpredictable Socioeconomic Costs,” *Journal of Economic Dynamics and Control*, April 2024, *161*, 104841.

**Benson, Christopher L. and Christopher L. Magee**, “Quantitative Determination of Technological Improvement from Patent Data,” *PLOS ONE*, April 2015, *10* (4), e0121635.

**Bento, Nuno and Charlie Wilson**, “Measuring the Duration of Formative Phases for Energy Technologies,” *Environmental Innovation and Societal Transitions*, December 2016, *21*, 95–112.

- Berckmans, Gert, Maarten Messagie, Jelle Smekens, Noshin Omar, Lieselot Vanhaverbeke, and Joeri Van Mierlo**, “Cost Projection of State of the Art Lithium-Ion Batteries for Electric Vehicles Up to 2030,” *Energies*, September 2017, *10* (9), 1314.
- Bhattacharyya, Subhes C. and Govinda R. Timilsina**, “A Review of Energy System Models,” *International Journal of Energy Sector Management*, November 2010, *4* (4), 494–518.
- Blondeel, Mathieu, James Price, Michael Bradshaw, Steve Pye, Paul Dodds, Caroline Kuzemko, and Gavin Bridge**, “Global Energy Scenarios: A Geopolitical Reality Check,” *Global Environmental Change*, January 2024, *84*, 102781.
- Bolinger, Mark, Ryan Wiser, and Eric O’Shaughnessy**, “Levelized Cost-Based Learning Analysis of Utility-Scale Wind and Solar in the United States,” *iScience*, June 2022, *25* (6), 104378.
- Boretos, George P.**, “The Suitability of S-Curves in Forecasting the Mobile Phone Market,” in “Mobile Telephones: Networks, Applications, and Performance,” Nova Publishers, 2008, pp. 253–264.
- Bowen, Alex and Cameron Hepburn**, “Green Growth: An Assessment,” *Oxford Review of Economic Policy*, October 2014, *30* (3), 407–422.
- Breiman, Leo**, “Heuristics of Instability and Stabilization in Model Selection,” *The Annals of Statistics*, December 1996, *24* (6), 2350–2383.
- Bretherton, Christopher S., Martin Widmann, Valentin P. Dymnikov, John M. Wallace, and Ileana Bladé**, “The Effective Number of Spatial Degrees of Freedom of a Time-Varying Field,” *Journal of Climate*, July 1999, *12*, 1990–2009.
- Breyer, Christian, Siavash Khalili, Dmitrii Bogdanov, Manish Ram, Ayobami Solomon Oyewo, Arman Aghahosseini, Ashish Gulagi, A A Solomon,**

Dominik Keiner, Gabriel Lopez, Poul Alberg Østergaard, Henrik Lund, Brian V Mathiesen, Mark Z Jacobson, Marta Victoria, Sven Teske, Thomas Pregger, Vasilis Fthenakis, Marco Raugei, Hannele Holttinen, Ugo Bardi, Auke Hoekstra, and Benjamin K Sovacool, “On the History and Future of 100% Renewable Energy Systems Research,” *IEEE Access*, July 2022, *10*, 78176–78218.

Bruno, Sergio, Shabbir Ahmed, Alexander Shapiro, and Alexandre Street, “Risk Neutral and Risk Averse Approaches to Multistage Renewable Investment Planning under Uncertainty,” *European Journal of Operational Research*, May 2016, *250* (3), 979–989.

Bukenberger, Jesse and Bryan Palmintier, “Stochastic Generation Capacity Expansion Planning with Approximate Dynamic Programming,” *2018 IEEE/PES Transmission and Distribution Conference and Exposition (T&D)*, April 2018, pp. 1–5.

Bunn, Derek W. and Ahti A. Salo, “Forecasting with Scenarios,” *European Journal of Operational Research*, August 1993, *68* (3), 291–303.

Byers, Edward, Volker Krey, Elmar Kriegler, Keywan Riahi, Roberto Schaeffer, Jarmo Kikstra, Robin Lamboll, Zebedee Nicholls, Marit Sandstad, Chris Smith, van der Wijst, Kaj, Al -Khourdajie, Alaa, Franck Lecocq, Portugal-Pereira, Joana, Yamina Saheb, Anders Stromman, Harald Winkler, Cornelia Auer, Elina Brutschin, Matthew Gidden, Philip Hackstock, Mathijs Harmen, Daniel Huppmann, Peter Kolp, Claire Lepault, Jared Lewis, Giacomo Marangoni, Müller-Casseres, Eduardo, Ragnhild Skeie, Michaela Werning, Katherine Calvin, Piers Forster, Celine Guivarch, Tomoko Hasegawa, Malte Meinshausen, Glen Peters, Joeri Rogelj, Bjorn Samset, Julia Steinberger, Massimo Tavoni, and van Vuuren, Detlef, “AR6 Scenarios Database,” <https://zenodo.org/record/5886911> November 2022.

- Bücker, Joris, R. Maria Del Rio-Chanona, Anton Pichler, Matthew C. Ives, and J. Doyme Farmer**, “Employment dynamics in a rapid decarbonization of the US power sector,” *Joule*, February 2025, *9* (2), 101803.
- Capellán-Pérez, Iñigo, Carlos de Castro, and Iñaki Arto**, “Assessing Vulnerabilities and Limits in the Transition to Renewable Energies: Land Requirements under 100% Solar Energy Scenarios,” *Renewable and Sustainable Energy Reviews*, September 2017, *77*, 760–782.
- Chai, Tianfeng and Roland Draxler**, “Root Mean Square Error (RMSE) or Mean Absolute Error (MAE)? – Arguments against Avoiding RMSE in the Literature,” *Geoscientific Model Development*, June 2014, *7* (3), 1247–1250.
- Chakravorty, Jhelum and Aditya Mahajan**, “Multi-Armed Bandits, Gittins Index, and Its Calculation,” in N. Balakrishnan, ed., *Methods and Applications of Statistics in Clinical Trials*, Hoboken, NJ, USA: John Wiley & Sons, Inc., June 2014, pp. 416–435.
- Chang, Nathan L., Bonna K. Newman, and Renate J. Egan**, “Future Cost Projections for Photovoltaic Module Manufacturing Using a Bottom-up Cost and Uncertainty Model,” *Solar Energy Materials and Solar Cells*, April 2022, *237*, 111529.
- Cherp, Aleh, Vadim Vinichenko, Jale Tosun, Joel A. Gordon, and Jessica Jewell**, “National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets,” *Nature Energy*, 2021, *6*, 742–754.
- Clements, Michael P. and David F. Hendry, eds**, *A Companion to Economic Forecasting* Blackwell Companions to Sociology, Hoboken: Wiley, 2008.
- Comin, Diego and Bart Hobijn**, “An Exploration of Technology Diffusion,” *The American Economic Review*, 2010, *100* (5), 2031–2059.
- , – , and **Emilie Rovito**, “Five Facts You Need to Know About Technology Diffusion,” *NBER Working Paper*, January 2006.

- Conejo, Antonio J. and Luis Baringo**, *Power System Operations Power Electronics and Power Systems*, Cham: Springer, 2018.
- Cowan, Robin**, “Tortoises and Hares: Choice Among Technologies of Unknown Merit,” *The Economic Journal*, 1991, *101* (407), 801–814.
- Creutzig, Felix, Peter Agoston, Jan Christoph Goldschmidt, Gunnar Luderer, Gregory F. Nemet, and Robert C. Pietzcker**, “The Underestimated Potential of Solar Energy to Mitigate Climate Change,” *Nature Energy*, August 2017, *2* (9).
- de Castro, Carlos, Margarita Mediavilla, Luis Javier Miguel, and Fernando Frechoso**, “Global Solar Electric Potential: A Review of Their Technical and Sustainable Limits,” *Renewable and Sustainable Energy Reviews*, December 2013, *28*, 824–835.
- Dietz, Simon and Cameron Hepburn**, “Benefit–Cost Analysis of Non-Marginal Climate and Energy Projects,” *Energy Economics*, November 2013, *40*, 61–71.
- Diffendorfer, Jay E., Brian Sergi, Anthony Lopez, Travis Williams, Michael Gleason, Zach Ancona, and Wesley Cole**, “The Interplay of Future Solar Energy, Land Cover Change, and Their Projected Impacts on Natural Lands and Croplands in the US,” *Science of The Total Environment*, October 2024, *947*, 173872.
- Dixit, Robert K. and Robert S. Pindyck**, *Investment under Uncertainty*, Princeton University Press, December 1994.
- Dooley, James J.**, “Estimating the Supply and Demand for Deep Geologic CO<sub>2</sub> Storage Capacity over the Course of the 21st Century: A Meta-analysis of the Literature,” *Energy Procedia*, January 2013, *37*, 5141–5150.
- Dowell, Niall Mac, Paul S. Fennell, Nilay Shah, and Geoffrey C. Maitland**, “The Role of CO<sub>2</sub> Capture and Utilization in Mitigating Climate Change,” *Nature Climate Change*, April 2017, *7* (4), 243–249.

- Ek, Kristina and Patrik Söderholm**, “Technology Learning in the Presence of Public R\&D: The Case of European Wind Power,” *Ecological Economics*, October 2010, *69* (12), 2356–2362.
- Elia, A., M. Taylor, B. Ó Gallachóir, and F. Rogan**, “Wind Turbine Cost Reduction: A Detailed Bottom-up Analysis of Innovation Drivers,” *Energy Policy*, December 2020, *147*, 111912.
- Elshurafa, Amro M., Shahad R. Albardi, Simona Bigerna, and Carlo Andrea Bollino**, “Estimating the Learning Curve of Solar PV Balance-of-System for over 20 Countries: Implications and Policy Recommendations,” *Journal of Cleaner Production*, September 2018, *196*, 122–134.
- Farmer, J. Doyne**, *Making Sense of Chaos: A Better Economics for a Better World*, Dublin: Allen Lane, 2024.
- **and François Lafond**, “How Predictable Is Technological Progress?,” *Research Policy*, April 2016, *45* (3), 647–665.
- , **Cameron Hepburn, Matthew C. Ives, Thomas Hale, Thom Wetzer, Penny Mealy, Ryan Rafaty, Sugandha Srivastav, and Rupert Way**, “Sensitive Intervention Points in the Post-Carbon Transition,” *Science*, April 2019, *364* (6436), 132–134.
- , – , **Penny Mealy, and Alexander Teytelboym**, “A Third Wave in the Economics of Climate Change,” *Climate Research in Gothenburg on*, 2015, *62*, 329–357.
- Farrow, R. Scott**, “A Missing Error Term in Benefit–Cost Analysis,” *ACS Publications*, February 2012.
- Feldman, David, Vignesh Ramasamy, Ran Fu, Ashwin Ramdas, Jal Desai, and Robert Margolis**, “U.S. Solar Photovoltaic System and Energy Storage Cost Benchmark

(Q1 2020),” Technical Report NREL/TP-6A20-77324, 1764908, MainId:26270 January 2021.

**Ferioli, Francesco and Bob van der Zwaan**, “Learning in Times of Change: A Dynamic Explanation for Technological Progress,” *Environmental Science and Technology*, June 2009, *43* (11), 4002–4008.

– , **Koen Schoots, and Bob van der Zwaan**, “Use and Limitations of Learning Curves for Energy Technology Policy: A Component-Learning Hypothesis,” *Energy Policy*, July 2009, *37* (7), 2525–2535.

**Fernandes, Bartolomeu, Jorge Cunha, and Paula Ferreira**, “The Use of Real Options Approach in Energy Sector Investments,” *Renewable and Sustainable Energy Reviews*, December 2011, *15* (9), 4491–4497.

**Fildes, Robert and Spyros Makridakis**, “The Impact of Empirical Accuracy Studies on Time Series Analysis and Forecasting,” *International Statistical Review / Revue Internationale de Statistique*, 1995, *63* (3), 289–308.

**Fisher, Jacob C. and Robert H. Pry**, “A Simple Substitution Model of Technological Change,” *Technological Forecasting and Social Change*, January 1971, *3*, 75–88.

**Fouquet, Roger, ed.**, *Handbook on Green Growth*, Cheltenham, UK ; Northampton, MA: Edward Elgar Publishing, 2019.

**Funk, Jeffrey L. and Christopher L. Magee**, “Rapid Improvements with No Commercial Production: How Do the Improvements Occur?,” *Research Policy*, April 2015, *44* (3), 777–788.

**Gallagher, Kelly Sims, Arnulf Grübler, Laura Kuhl, Gregory F. Nemet, and Charlie Wilson**, “The Energy Technology Innovation System,” *Annual Review of Environment and Resources*, November 2012, *37* (1), 137–162.

- Gittins, John C.**, “Bandit Processes and Dynamic Allocation Indices,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1979, 41 (2), 148–177.
- , **Richard Weber, and Kevin D. Glazebrook**, *Multi-Armed Bandit Allocation Indices*, second ed., Hoboken, NJ: John Wiley & Sons, 2011.
- Glenk, Gunther, Rebecca Meier, and Stefan Reichelstein**, “Cost Dynamics of Clean Energy Technologies,” *Schmalenbach Journal of Business Research*, 2021, 73, 179–206.
- Gluzberg, Victor E. and Yuri A. Katz**, “Probabilistic Forecasting of Transition towards Sustainable Energy Production,” *Journal of Climate Finance*, March 2025, 10, 100063.
- Gneiting, Tilmann and Adrian E Raftery**, “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association*, March 2007, 102 (477), 359–378.
- and **Matthias Katzfuss**, “Probabilistic Forecasting,” *Annual Review of Statistics and Its Application*, January 2014, 1 (1), 125–151.
- , **Fadoua Balabdaoui, and Adrian E. Raftery**, “Probabilistic Forecasts, Calibration and Sharpness,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, April 2007, 69 (2), 243–268.
- Grafström, Jonas and Rahmatallah Poudineh**, “A Critical Assessment of Learning Curves for Solar and Wind Power Technologies,” Technical Report 43, The Oxford Institute for Energy Studies, Oxford 2021.
- Grant, Neil, Ajay Gambhir, Shivika Mittal, Chris Greig, and Alexandre C. Köberle**, “Enhancing the Realism of Decarbonisation Scenarios with Practicable Regional Constraints on CO<sub>2</sub> Storage Capacity,” *International Journal of Greenhouse Gas Control*, October 2022, 120, 103766.

Grubb, Michael, Paul Drummond, Jean-Francois Mercure, Cameron Hepburn, Peter Barbrook-Johnson, Carlos Fjoao Joao Carlos Ferraz, Alex Clark, Laura Diaz Anadon, J. Doyne Farmer, Ben Hinder, Matthew C. Ives, Aled Jones, Gao Jun, Ulka Kelkar, Sergey Kolesnikov, Aileen Lam, Ritu Mathur, Roberto Pasqualino, Cristina Penasco, Hector Pollitt, Luma Ramos, Andrea Roventini, Pablo Salas, Simon Sharpe, Zhu Songli, Pim Vercoulen, Kamna Waghray, and Zhang Xiliang, “The New Economics of Innovation and Transition: Evaluating Opportunities and Risks,” Technical Report, EEIST 2021.

Gruber, Harald, “Learning by Doing and Spillovers: Further Evidence for the Semiconductor Industry,” *Review of Industrial Organization*, 1998, *13*, 697–711.

Grübler, Arnulf, *The Rise and Fall of Infrastructures: Dynamics of Evolution and Technological Change in Transport*, Heidelberg: Physica-Verlag, 1990.

– , “Diffusion Long-Term Patterns and Discontinuities,” *Technological Forecasting and Social Change*, 1991, *39*, 159–180.

– and Nebojša Nakićenović, “Long Waves, Technology Diffusion, and Substitution,” *Review (Fernand Braudel Center)*, 1991, *14* (2), 313–343.

– and – , “Decarbonizing the Global Energy System,” *Technological Forecasting and Social Change*, September 1996, *53* (1), 97–110.

– , Charlie Wilson, and Gregory F. Nemet, “Apples, Oranges, and Consistent Comparisons of the Temporal Dynamics of Energy Transitions,” *Energy Research & Social Science*, December 2016, *22*, 18–25.

– , Nebojša Nakićenović, and David G. Victor, “Dynamics of Energy Technologies and Global Change,” *Energy Policy*, May 1999, *27* (5), 247–280.

Grunfeld, Yehuda and Zvi Griliches, “Is Aggregation Necessarily Bad?,” *The Review of Economics and Statistics*, 1960, *42* (1), 1–13.

- Hach, Daniel and Stefan Spinler**, “Robustness of Capacity Markets: A Stochastic Dynamic Capacity Investment Model,” *OR Spectrum*, March 2018, *40* (2), 517–540.
- Hansen, Jan P., Patrick A. Narbel, and Dag L. Aksnes**, “Limits to Growth in the Renewable Energy Sector,” *Renewable and Sustainable Energy Reviews*, April 2017, *70*, 769–774.
- Harberger, Arnold C.**, “Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay,” *Journal of Economic Literature*, 1971, *9* (3), 785–797.
- Hayashi, Daisuke, Joern Huenteler, and Joanna I. Lewis**, “Gone with the Wind: A Learning Curve Analysis of China’s Wind Power Industry,” *Energy Policy*, September 2018, *120*, 38–51.
- Hayward, Jennifer A. and Paul W. Graham**, “A Global and Local Endogenous Experience Curve Model for Projecting Future Uptake and Cost of Electricity Generation Technologies,” *Energy Economics*, November 2013, *40*, 537–548.
- Heij, Christiaan, ed.**, *Econometric Methods with Applications in Business and Economics*, 1. publ ed., Oxford New York: Oxford University Press, 2004.
- Hendry, David F. and Kirstin Hubrich**, “Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate,” *Journal of Business & Economic Statistics*, 2011, *29* (2), 216–227.
- Herbst, Andrea, Felipe Toro, Felix Reitze, and Eberhard Jochem**, “Introduction to Energy Systems Modelling,” *Swiss Journal of Economics and Statistics*, April 2012, *148* (2), 111–135.
- Hernandez, Rebecca R., Madison K. Hoffacker, and Christopher B. Field**, “Efficient Use of Land to Meet Sustainable Energy Needs,” *Nature Climate Change*, April 2015, *5* (4), 353–358.

- Heuberger, Clara F., Edward S. Rubin, Iain Staffell, Nilay Shah, and Niall Mac Dowell**, “Power Capacity Expansion Planning Considering Endogenous Technology Cost Learning,” *Applied Energy*, 2017, *204*, 831–845.
- Hoffman, Kenneth C. and David O. Wood**, “Energy System Modeling and Forecasting,” *Annual Review of Environment and Resources*, November 1976, *1* (Volume 1, 1976), 423–453.
- Höök, Mikael, Junchen Li, Kersti Johansson, and Simon Snowden**, “Growth Rates of Global Energy Systems and Future Outlooks,” *Natural Resources Research*, March 2012, *21* (1), 23–41.
- Hötte, Kerstin and Su Jung Jee**, “Knowledge for a Warmer World: A Patent Analysis of Climate Change Adaptation Technologies,” *Technological Forecasting and Social Change*, October 2022, *183*, 121879.
- Hsieh, I-Yun Lisa, Menghsuan Sam Pan, Yet-Ming Chiang, and William H. Green**, “Learning Only Buys You so Much: Practical Limits on Battery Price Reduction,” *Applied Energy*, April 2019, *239*, 218–224.
- Ingersoll, Jonathan E.**, *Theory of Financial Decision Making* Rowman & Littlefield Studies in Financial Economics, third (print) ed., Savage, MD: Rowman & Littlefield, 1987.
- IPCC**, *Climate Change 2022, Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK and New York, NY, USA: IPCC, 2022.
- IPCC**, *Climate Change 2022, Impacts, Adaptation and Vulnerability: Working Group II Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 1 ed., Cambridge University Press, June 2023.

- **and WMO, eds**, *Climate Change: The 1990 and 1992 IPCC Assessments, IPCC First Assessment Report Overview and Policymaker Summaries and 1992 IPCC Supplement*, Geneva: IPCC, 1992.
- , **ed.**, *Climate Change 2007, Mitigation of Climate Change: Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge ; New York: Cambridge University Press, 2007.
- Irwin, Douglas A. and Peter J. Klenow**, “Learning-by-Doing Spillovers in the Semiconductor Industry,” *Source: Journal of Political Economy*, 1994, *102* (6), 1200–1227.
- Iyer, Gokul, Nathan Hultman, Jiyong Eom, Haewon McJeon, Pralit Patel, and Leon Clarke**, “Diffusion of Low-Carbon Technologies and the Feasibility of Long-Term Climate Targets,” *Technological Forecasting and Social Change*, January 2015, *90*, 103–118.
- Jebaraj, Sargunam and Sellvarasan K. Iniyan**, “A Review of Energy Models,” *Renewable and Sustainable Energy Reviews*, August 2006, *10* (4), 281–311.
- Jonson, Emma, Christian Azar, Kristian Lindgren, and Liv Lundberg**, “Exploring the Competition between Variable Renewable Electricity and a Carbon-Neutral Baseload Technology,” *Energy Systems*, February 2020, *11* (1), 21–44.
- Junginger, Martin, André Faaij, and Wim C Turkenburg**, “Cost Reduction Prospects for Offshore Wind Farms,” *Wind Engineering*, 2004, *28* (1), 97–118.
- Kavлак, Goksin, James McNerney, and Jessika E. Trancik**, “Evaluating the Causes of Cost Reduction in Photovoltaic Modules,” *Energy Policy*, December 2018, *123*, 700–710.
- Kazlou, Tsimafei, Aleh Cherp, and Jessica Jewell**, “Feasible Deployment of Carbon Capture and Storage and the Requirements of Climate Targets,” *Nature Climate Change*, October 2024, *14* (10), 1047–1055.

**Keppo, Ilkka and Manfred Strubegger**, “Short Term Decisions for Long Term Problems – The Effect of Foresight on Model Based Energy Systems Analysis,” *Energy*, May 2010, *35* (5), 2033–2042.

**Klemun, Magdalena M., Goksin Kavlak, James McNerney, and Jessika E. Trancik**, “Mechanisms of Hardware and Soft Technology Evolution and the Implications for Solar Energy Cost Trends,” *Nature Energy*, August 2023, *8* (8), 827–838.

**Kothari, Sumit, Nadia Ameli, and Michael Grubb**, “Drivers of Balance-of-System Costs for Solar PV Technology,” 2023.

**Krey, Volker**, “Global Energy-Climate Scenarios and Models: A Review,” *Wiley Interdisciplinary Reviews: Energy and Environment*, 2014, *3* (4), 363–383.

– , **Fei Guo, Peter Kolp, Wenji Zhou, Roberto Schaeffer, Aayushi Awasthy, Christoph Bertram, Harmen Sytze de Boer, Panagiotis Fragkos, Shinichiro Fujimori, Chenmin He, Gokul Iyer, Kimon Keramidas, Alexandre C. Köberle, Ken Oshiro, Lara Aleluia Reis, Bianka Shoai-Tehrani, Saritha Vishwanathan, Pantelis Capros, Laurent Drouet, James E. Edmonds, Amit Garg, David E.H.J. Gernaat, Kejun Jiang, Maria Kannavou, Alban Kitous, Elmar Kriegler, Gunnar Luderer, Ritu Mathur, Matteo Muratori, Fuminori Sano, and Detlef P. van Vuuren**, “Looking under the Hood: A Comparison of Techno-Economic Assumptions across National and Global Integrated Assessment Models,” *Energy*, April 2019, *172*, 1254–1267.

**Kumbaroğlu, Gürkan, Reinhard Madlener, and Mustafa Demirel**, “A Real Options Evaluation Model for the Diffusion Prospects of New Renewable Power Generation Technologies,” *Energy Economics*, July 2008, *30* (4), 1882–1908.

**Labriet, Maryse, Claire Nicolas, Stéphane Tchung-Ming, Amit Kanudia, and Richard Loulou**, “Energy Decisions in an Uncertain Climate and Technology Outlook:

How Stochastic and Robust Methodologies Can Assist Policy-Makers,” in George Gianakidis, Maryse Labriet, Brian Ó Gallachóir, and GianCarlo Tosato, eds., *Informing Energy and Climate Policies Using Energy Systems Models: Insights from Scenario Analysis Increasing the Evidence Base*, Lecture Notes in Energy, Cham: Springer International Publishing, 2015, pp. 69–91.

**Lafond, François**, “Forecasting Technological Progress,” March 2025.

– , **Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J. Doyne Farmer**, “How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts,” *Technological Forecasting and Social Change*, March 2018, *128*, 104–117.

– , **Diana Seave Greenwald, and J. Doyne Farmer**, “Can Stimulating Demand Drive Costs Down? World War II as a Natural Experiment,” *SSRN Electronic Journal*, June 2020.

**Lane, Joe, Chris Greig, and Andrew Garnett**, “Uncertain Storage Prospects Create a Conundrum for Carbon Capture and Storage Ambitions,” *Nature Climate Change*, November 2021, *11* (11), 925–936.

**Leggett, Jeremy, William J. Pepper, and Robert J. Swart**, “Emissions Scenarios for the IPCC: An Update,” 1992.

**Lehtila, Antti and Richard Loulou**, “TIMES Damage Functions,” November 2005.

**Levin, Nissan, Asher Tishler, and Jacob Zahavi**, “Time Step vs. Dynamic Optimization of Generation-Capacity-Expansion Programs of Power Systems,” *Operations Research*, 1983, *31* (5), 891–914.

**Link, Steffen, Lara Schneider, Annegret Stephan, Lukas Weymann, and Patrick Plötz**, “Feasibility of Meeting Future Battery Demand via Domestic Cell Production in Europe,” *Nature Energy*, March 2025, pp. 1–9.

- Lopion, Peter, Peter Markewitz, Detlef Stolten, and Martin Robinius**, “Cost Uncertainties in Energy System Optimization Models: A Quadratic Programming Approach for Avoiding Penny Switching Effects,” *Energies*, 2019, 12.
- , – , **Martin Robinius, and Detlef Stolten**, “A Review of Current Challenges and Trends in Energy Systems Modeling,” *Renewable and Sustainable Energy Reviews*, November 2018, 96, 156–166.
- Louwen, Atse, Steffi Schreiber, and Martin Junginger**, “Implementation of Experience Curves in Energy-System Models,” *Technological Learning in the Transition to a Low-Carbon Energy System: Conceptual Issues, Empirical Findings, and Use, in Energy Modeling*, January 2019, pp. 33–47.
- Lovins, Amory B.**, “Energy Strategy: The Road Not Taken?,” *New York Times*, November 1977, (20).
- Madsen, Dorte Nørgaard and Jan Petter Hansen**, “Outlook of Solar Energy in Europe Based on Economic Growth Characteristics,” *Renewable and Sustainable Energy Reviews*, October 2019, 114, 109306.
- Marinakis, Yorgos D.**, “Forecasting Technology Diffusion with the Richards Model,” *Technological Forecasting and Social Change*, January 2012, 79 (1), 172–179.
- Markowitz, Harry**, “Portfolio Selection,” *The Journal of Finance*, 1952, 7 (1), 77–91.
- , “Mean–Variance Approximations to Expected Utility,” *European Journal of Operational Research*, April 2014, 234 (2), 346–355.
- McLaren, Duncan and Nils Markusson**, “The Co-Evolution of Technological Promises, Modelling, Policies and Climate Change Targets,” *Nature Climate Change*, May 2020, 10 (5), 392–397.

**McNerney, James, J. Doayne Farmer, and Jessika E. Trancik**, “Historical Costs of Coal-Fired Electricity and Implications for the Future,” *Energy Policy*, June 2011, *39* (6), 3042–3054.

– , – , **Sidney Redner, and Jessika E. Trancik**, “Role of Design Complexity in Technology Improvement,” *Proceedings of the National Academy of Sciences*, May 2011, *108* (22), 9008–9013.

**Mealy, Penny and Alexander Teytelboym**, “Economic Complexity and the Green Economy,” *Research Policy*, October 2022, *51* (8), 103948.

– **and Cameron Hepburn**, “Chapter 17: Transformational Change: Parallels for Addressing Climate and Development Goals,” in “in” June 2020, chapter Handbook on the Economics of Climate Change.

– , **Pete Barbrook-Johnson, Matthew C. Ives, Sugandha Srivastav, and Cameron Hepburn**, “Sensitive Intervention Points: A Strategic Approach to Climate Action,” *Oxford Review of Economic Policy*, November 2023, *39* (4), 694–710.

**Meng, Jing, Rupert Way, Elena Verdolini, and Laura Diaz Anadon**, “Comparing Expert Elicitation and Model-Based Probabilistic Technology Cost Forecasts for the Energy Transition,” *Proceedings of the National Academy of Sciences of the United States of America*, July 2021, *118* (27).

**Mercure, Jean-Francois, Hector Pollitt, Neil R. Edwards, Philip B. Holden, Unnada Chewprecha, Pablo Salas, Aileen Lam, Florian Knobloch, and Jorge E. Vinuales**, “Environmental Impact Assessment for Climate Change Policy with the Simulation-Based Integrated Assessment Model E3ME-FTT-GENIE,” *Energy Strategy Reviews*, April 2018, *20*, 195–208.

**Messner, Sabine, Alexander N. Golodnikov, and Andrii Gritsevskii**, “A Stochastic Version of the Dynamic Linear Programming Model MESSAGE III,” *Energy*, September 1996, *21* (9), 775–784.

- Milford, James, Max Henrion, Chad Hunter, Emily Newes, Caroline Hughes, and Samuel F. Baldwin**, “Energy Sector Portfolio Analysis with Uncertainty,” *Applied Energy*, January 2022, *306*, 117926.
- Moore, Gordon E.**, “Progress in Digital Integrated Electronics,” *IEEE*, 1975, pp. 11–13.
- Murphy, Allan H. and Robert L. Winkler**, “A General Framework for Forecast Verification,” *Monthly Weather Review*, July 1987, *115*, 1330–1338.
- Nagy, Bela, J. Doayne Farmer, Quan M. Bui, and Jessika E. Trancik**, “Statistical Basis for Predicting Technological Progress,” *PLoS ONE*, February 2013, *8* (2), 52669.
- Napp, Tamaryn, Dan Bernie, Rebecca Thomas, Jason Lowe, Adam Hawkes, and Ajay Gambhir**, “Exploring the Feasibility of Low-Carbon Scenarios Using Historical Energy Transitions Analysis,” *Energies*, January 2017, *10* (1), 116.
- Neij, Lena**, “Cost Development of Future Technologies for Power Generation—A Study Based on Experience Curves and Complementary Bottom-up Assessments,” *Energy Policy*, June 2008, *36* (6), 2200–2211.
- **and Gregory F. Nemet**, “Accelerating the Low-Carbon Transition Will Require Policy to Enhance Local Learning,” *Energy Policy*, August 2022, *167*, 113043.
- Nemet, Gregory F.**, “Beyond the Learning Curve: Factors Influencing Cost Reductions in Photovoltaics,” *Energy Policy*, November 2006, *34* (17), 3218–3232.
- , “Demand-Pull, Technology-Push, and Government-Led Incentives for Non-Incremental Technical Change,” *Research Policy*, June 2009, *38* (5), 700–709.
- , “Inter-Technology Knowledge Spillovers for Energy Technologies,” *Energy Economics*, September 2012, *34* (5), 1259–1270.
- **and Evan Johnson**, “Do Important Inventions Benefit from Knowledge Originating in Other Technological Domains?,” *Research Policy*, February 2012, *41* (1), 190–200.

- , **Jenna Greene, Finn Müller-Hansen, and Jan C. Minx**, “Dataset on the Adoption of Historical Technologies Informs the Scale-up of Emerging Carbon Dioxide Removal Measures | Communications Earth & Environment,” *Communications Earth & Environment*, October 2023, *4* (1), 397.
- , **Jiaqi Lu, Varun Rai, and Rohan Rao**, “Knowledge Spillovers between PV Installers Can Reduce the Cost of Installing Solar PV,” *Energy Policy*, September 2020, *144*, 111600.
- Nijnens, Joey, Paul Behrens, Oscar Kraan, Benjamin Sprecher, and René Kleijn**, “Energy Transition Will Require Substantially Less Mining than the Current Fossil System,” *Joule*, November 2023, *7* (11), 2408–2413.
- Nonhebel, Sanderine**, “Renewable Energy and Food Supply: Will There Be Enough Land?,” *Renewable and Sustainable Energy Reviews*, April 2005, *9* (2), 191–201.
- Nordhaus, William D.**, “An Optimal Transition Path for Controlling Greenhouse Gases,” *Science*, 1992, *258* (5086), 1315–1319.
- , “The Perils of the Learning Model for Modeling Endogenous Technological Change,” *Energy Journal*, January 2014, *35* (1), 1–13.
- Odenweller, Adrian and Falko Ueckerdt**, “The Green Hydrogen Ambition and Implementation Gap,” *Nature Energy*, January 2025, *10*, 110–123.
- , – , **Gregory F. Nemet, Miha Jensterle, and Gunnar Luderer**, “Probabilistic Feasibility Space of Scaling up Green Hydrogen Supply,” *Nature Energy*, September 2022, *7*, 854–865.
- Olivetti, Elsa A., Gerbrand Ceder, Gabrielle G. Gaustad, and Xinkai Fu**, “Lithium-Ion Battery Supply Chain Considerations: Analysis of Potential Bottlenecks in Critical Metals,” *Joule*, October 2017, *1* (2), 229–243.

- Pant, P. Narayan and William H. Starbuck, “Innocents in the Forest: Forecasting and Research Methods,” *Journal of Management*, June 1990, 16 (2), 433–460.
- Pesaran, M. Hashem, Andreas Pick, and Allan Timmermann, “Forecasting with Panel Data: Estimation Uncertainty versus Parameter Heterogeneity,” April 2024.
- Peskir, Goran and Albert Shiryaev, *Optimal Stopping and Free-Boundary Problems* Lectures in Mathematics ETH Zürich, Basel ; Boston: Birkhäuser Verlag, 2006.
- Petropoulos, Fotios, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K. Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J. Bessa, Jakub Bijak, John E. Boylan, Jethro Browell, Claudio Carnevale, Jennifer L. Castle, Pasquale Cirillo, Michael P. Clements, Clara Cordeiro, Fernando Luiz Cyrino Oliveira, Shari De Baets, Alexander Dokumentov, Joanne Ellison, Piotr Fiszeder, Philip Hans Franses, David T. Frazier, Michael Gilliland, M. Sinan Gönül, Paul Goodwin, Luigi Grossi, Yael Grushka-Cockayne, Mariangela Guidolin, Massimo Guidolin, Ulrich Gunter, Xiaojia Guo, Renato Guseo, Nigel Harvey, David F. Hendry, Ross Hollyman, Tim Januschowski, Jooyoung Jeon, Victor Richmond R. Jose, Yanfei Kang, Anne B. Koehler, Stephan Kolassa, Nikolaos Kourentzes, Sonia Leva, Feng Li, Konstantia Litsiou, Spyros Makridakis, Gael M. Martin, Andrew B. Martinez, Sheik Meeran, Theodore Modis, Konstantinos Nikolopoulos, Dilek Önköl, Alessia Paccagnini, Anastasios Panagiotelis, Ioannis Panapakidis, Jose M. Pavía, Manuela Pedio, Diego J. Pedregal, Pierre Pinson, Patrícia Ramos, David E. Rapach, J. James Reade, Bahman Rostami-Tabar, Michał Rubaszek, Georgios Sermpinis, Han Lin Shang, Evangelos Spiliotis, Aris A. Syntetos, Priyanga Dilini Talagala, Thiyanga S. Talagala, Len Tashman, Dimitrios Thomakos, Thordis Thorarinsdottir, Ezio Todini, Juan Ramón Trapero Arenas, Xiaoqian Wang, Robert L. Winkler, Alisa Yusupova, and Florian

- Ziel**, “Forecasting: Theory and Practice,” *International Journal of Forecasting*, July 2022, 38 (3), 705–871.
- Pfenninger, Stefan, Adam Hawkes, and James Keirstead**, “Energy Systems Modeling for Twenty-First Century Energy Challenges,” *Renewable and Sustainable Energy Reviews*, 2014, 33, 74–86.
- Pichler, Anton, Francois Lafond, and J. Doyne Farmer**, “Technological Interdependencies Predict Innovation Dynamics,” *SSRN Electronic Journal*, 2020.
- Pielke, Roger and Justin Ritchie**, “Distorting the View of Our Climate Future: The Misuse and Abuse of Climate Pathways and Scenarios,” *Energy Research and Social Science*, February 2021, 72.
- Pindyck, Robert S.**, “Climate Change Policy: What Do the Models Tell Us?,” *NBER Working Paper*, July 2013.
- Powell, Warren B.**, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Hoboken, US: John Wiley & Sons, Incorporated, 2011.
- Prina, Matteo Giacomo, Giampaolo Manzolini, David Moser, Benedetto Nastasi, and Wolfram Sparber**, “Classification and Challenges of Bottom-up Energy System Models - A Review,” *Renewable and Sustainable Energy Reviews*, September 2020, 129 (109917).
- Rai, Varun, Ariane Beck, Gregory F. Nemet, Xue Gao, Douglas Hannah, Adam Henry, and Erik Funkhouser**, “Knowledge Spillovers and Cost Reductions in Solar Soft Costs,” Technical Report DOE-UT-EE0007658, Univ. of Texas, Austin, TX (United States) October 2020.
- Ramsey, Frank P.**, “A Mathematical Theory of Saving,” *The Economic Journal*, 1928, 38 (152), 543–559.

- Raupach-Sumiya, Jörg, Hironao Matsubara, Andreas Prah, Astrid Aretz, and Steven Salecki**, “Regional Economic Effects of Renewable Energies - Comparing Germany and Japan,” *Energy, Sustainability and Society*, December 2015, 5 (1), 10.
- Reiner, David M.**, “Learning through a Portfolio of Carbon Capture and Storage Demonstration Projects,” *Nature Energy*, January 2016, 1 (1), 1–7.
- Ritchie, Hannah**, “How Does the Land Use of Different Electricity Sources Compare?,” *Our World in Data*, June 2022.
- Riva, Alberto Dalla, Janos Hethey, Silke Luers, Anna-Kathrin Wallasch, Knud Rehfeldt, Aidan Duffy, David Edward Weir, Maria Stenkvis, Andreas Uihlein, Tyler Stehly, and Eric Lantz**, “IEA Wind TCP Task 26: Wind Technology, Cost, and Performance Trends in Denmark, Germany, Ireland, Norway, Sweden, the European Union, and the United States: 2008-2016,” Technical Report, National Renewable Energy Laboratory (NREL), Golden, CO (US) November 2018.
- Rodrik, Dani**, “Green Industrial Policy,” *Oxford Review of Economic Policy*, 2014, 30 (3), 469–491.
- Rosenow, Samuel Kaspar and Penny Mealy**, “Turning Risks into Reward: Diversifying the Global Value Chains of Decarbonization Technologies,” *World Bank Policy Research Working Paper*, February 2024, (10696).
- Rubin, Edward S., Inês M.L. Azevedo, Paulina Jaramillo, and Sonia Yeh**, “A Review of Learning Rates for Electricity Supply Technologies,” *Energy Policy*, November 2015, 86, 198–218.
- Rudloff, Birgit, Alexandre Street, and Davi M. Valladão**, “Time Consistency and Risk Averse Dynamic Decision Models: Definition, Interpretation and Practical Consequences,” *European Journal of Operational Research*, May 2014, 234 (3), 743–750.

- Sahal, Devendra**, “A Theory of Progress Functions,” *A I I E Transactions*, March 1979, *11* (1), 23–29.
- Salo, Ahti, Michalis Doumpos, Juuso Liesiö, and Constantin Zopounidis**, “Fifty Years of Portfolio Optimization,” *European Journal of Operational Research*, October 2024, *318* (1), 1–18.
- Samadi, Sascha**, “The Experience Curve Theory and Its Application in the Field of Electricity Generation Technologies – A Literature Review,” *Renewable and Sustainable Energy Reviews*, February 2018, *82*, 2346–2364.
- Samuelson, Paul A.**, “Lifetime Portfolio Selection By Dynamic Stochastic Programming,” *The Review of Economics and Statistics*, 1969, *51* (3), 239–246.
- Schmidt, Oliver, Adam Hawkes, Ajay Gambhir, and Iain Staaell**, “The Future Cost of Electrical Energy Storage Based on Experience Rates,” *Nature Energy*, 2017, *2* (17110).
- Scott, Steven L.**, “A Modern Bayesian Look at the Multi-Armed Bandit,” *Applied Stochastic Models in Business and Industry*, 2010, *26* (6), 639–658.
- Semieniuk, Gregor, Philip B. Holden, Jean-Francois Mercure, Pablo Salas, Hector Pollitt, Katharine Jobson, Pim Vercoulen, Unnada Chewpreecha, Neil R. Edwards, and Jorge E. Viñuales**, “Stranded Fossil-Fuel Assets Translate to Major Losses for Investors in Advanced Economies,” *Nature Climate Change*, June 2022, *12* (6), 532–538.
- Shmueli, Galit**, “To Explain or to Predict?,” *Statistical Science*, August 2010, *25* (3), 289–310.
- Sievert, Katrin, Tobias S. Schmidt, and Bjarne Steffen**, “Considering Technology Characteristics to Project Future Costs of Direct Air Capture,” *Joule*, April 2024, *8* (4), 979–999.

- Sinha, Aditya, Aranya Venkatesh, Katherine Jordan, Cameron Wade, Hadi Eshraghi, Anderson R. De Queiroz, Paulina Jaramillo, and Jeremiah X. Johnson**, “Diverse decarbonization pathways under near cost-optimal futures,” *Nature Communications*, 2024, 15 (1), 8165.
- Sinsel, Simon R., Jochen Markard, and Volker H. Hoffmann**, “How Deployment Policies Affect Innovation in Complementary Technologies—Evidence from the German Energy Transition,” *Technological Forecasting and Social Change*, December 2020, 161.
- Smart, Bridget and J. Doyne Farmer**, “Macro and Micro Prediction Models,” October 2024.
- Smulders, Sjak, Michael Toman, and Cees Withagen**, “Growth Theory and ‘Green Growth’,” *Oxford Review of Economic Policy*, October 2014, 30 (3), 423–446.
- Söderholm, Patrik and Thomas Sundqvist**, “Empirical Challenges in the Use of Learning Curves for Assessing the Economic Prospects of Renewable Energy Technologies,” *Renewable Energy*, December 2007, 32 (15), 2559–2578.
- Sørensen, Bent**, “Energy and Resources,” *Science*, July 1975, 189 (4199), 255–260.
- Srivastav, Sugandha**, “Lost Potential: How the Kidnapping of a Solar Energy Pioneer Impacted the Cost of Renewable Energy and the Climate Crisis,” October 2023.
- Starrett, David A.**, *Foundations in Public Economics* Cambridge Economic Handbooks, Cambridge: Cambridge University Press, 1988.
- Steffen, Bjarne**, “The Importance of Project Finance for Renewable Energy Projects,” *Energy Economics*, January 2018, 69, 280–294.
- , “Estimating the Cost of Capital for Renewable Energy Projects,” *Energy Economics*, May 2020, 88.

- , **Martin Beuse, Paul Tautorat, and Tobias S. Schmidt**, “Experience Curves for Operations and Maintenance Costs of Renewable Energy Technologies,” *Joule*, February 2020, *4* (2), 359–375.
- Stephenson, David B. and Ian T. Jolliffe**, eds, *Forecast Verification: A Practitioner’s Guide in Atmospheric Science*, second ed., Hoboken: John Wiley & Sons, 2012.
- Stern, Nicholas and Chunping Xie**, “China’s New Growth Story: Linking the 14<sup>th</sup> Five-Year Plan with the 2060 Carbon Neutrality Pledge,” *Journal of Chinese Economic and Business Studies*, January 2023, *21* (1), 5–25.
- Strachan, Neil, Birgit Fais, and Hannah Daly**, “Reinventing the Energy Modelling–Policy Interface,” *Nature Energy*, February 2016, *1* (3), 1–3.
- Strupeit, Lars and Lena Neij**, “Cost Dynamics in the Deployment of Photovoltaics: Insights from the German Market for Building-Sited Systems,” *Renewable and Sustainable Energy Reviews*, March 2017, *69*, 948–960.
- Sutton, Richard S. and Andrew G. Barto**, *Reinforcement Learning: An Introduction*, Cambridge, Massachusetts: The MIT Press, 2014.
- Tàbara, David J., Niki Frantzeskaki, Katharina Hölscher, Simona Pedde, Kasper Kok, Francesco Lamperti, Jens H. Christensen, Jill Jäger, and Pam Berry**, “Positive Tipping Points in a Rapidly Warming World,” *Current Opinion in Environmental Sustainability*, April 2018, *31*, 120–129.
- Tashman, Leonard J.**, “Out-of-Sample Tests of Forecasting Accuracy: An Analysis and Review,” *International Journal of Forecasting*, October 2000, *16* (4), 437–450.
- Trancik, Jessika E., Goksin Kavlak, Magdalena M. Klemun, and Ajinkya S. Kamat**, “Modeling Photovoltaics Innovation and Deployment Dynamics,” Technical Report DOE-MIT-07662, Massachusetts Inst. of Technology (MIT), Cambridge, MA December 2020.

- Triulzi, Giorgio, Jeff Alstott, and Christopher L. Magee**, “Estimating Technology Performance Improvement Rates by Mining Patent Data,” *Technological Forecasting and Social Change*, September 2020, *158*, 120100.
- Trutnevyte, Evelina**, “Does Cost Optimization Approximate the Real-World Energy Transition?,” *Energy*, July 2016, *106*, 182–193.
- Turkovska, Olga, Katharina Gruber, Michael Klingler, Claude Klöckl, Luis Ramirez Camargo, Peter Regner, Sebastian Wehrle, and Johannes Schmidt**, “Methodological and Reporting Inconsistencies in Land-Use Requirements Misguide Future Renewable Energy Planning,” *One Earth*, October 2024, *7* (10), 1741–1759.
- Valero, Alicia, Antonio Valero, Guiomar Calvo, and Abel Ortego**, “Material Bottlenecks in the Future Development of Green Technologies,” *Renewable and Sustainable Energy Reviews*, October 2018, *93*, 178–200.
- van Beek, Lisette, Maarten Hajer, Peter Pelzer, Detlef van Vuuren, and Christophe Cassen**, “Anticipating Futures through Models: The Rise of Integrated Assessment Modelling in the Climate Science-Policy Interface since 1970,” *Global Environmental Change*, November 2020, *65*, 102191.
- van de Ven, Dirk-Jan, Iñigo Capellan-Peréz, Iñaki Arto, Ignacio Cazcarro, Carlos de Castro, Pralit Patel, and Mikel Gonzalez-Eguino**, “The Potential Land Requirements and Related Land Use Change Emissions of Solar Energy,” *Scientific Reports*, February 2021, *11* (1), 2907.
- van Sluisveld, Mariësse A.E., J.H.M. Harmsen, Nico Bauer, David L. McCollum, Keywan Riahi, Massimo Tavoni, Detlef P. van Vuuren, Charlie Wilson, and Bob van der Zwaan**, “Comparing Future Patterns of Energy System Change in 2°C Scenarios with Historically Observed Rates of Change,” *Global Environmental Change*, November 2015, *35*, 436–449.

- Vartiainen, Eero, Gaëtan Masson, Christian Breyer, David Moser, and Eduardo Román Medina**, “Impact of Weighted Average Cost of Capital, Capital Expenditure, and Other Parameters on Future Utility-Scale PV Levelised Cost of Electricity,” *Progress in Photovoltaics: Research and Applications*, 2020, 28 (6), 439–453.
- Vespucci, Maria Teresa, Marida Bertocchi, Paolo Pisciella, and Stefano Zigrino**, “Two-Stage Stochastic Mixed Integer Optimization Models for Power Generation Capacity Expansion with Risk Measures,” *Optimization Methods and Software*, March 2016, 31 (2), 305–327.
- Wagenvoort, Benjamin, Joel Dyer, François Lafond, and J. Doyne Farmer**, “The Universality of Technology Production,” 2025.
- Wagner, Fabian**, *Renewables in Future Power Systems Implications of Technological Learning and Uncertainty Renewables in Future Power Systems*, Springer International Publishing Switzerland, 2014.
- Wang, Weiguan and Johannes Ruf**, “Information Leakage in Backtesting,” June 2022.
- Way, Rupert, François Lafond, Fabrizio Lillo, Valentyn Panchenko, and J. Doyne Farmer**, “Wright Meets Markowitz: How Standard Portfolio Theory Changes When Assets Are Technologies Following Experience Curves,” *Journal of Economic Dynamics and Control*, April 2019, 101, 211–238.
- , **Matthew C. Ives, Penny Mealy, and J. Doyne Farmer**, “Empirically Grounded Technology Forecasts and the Energy Transition,” *Joule*, September 2022, 6, 1–26.
- Webster, Mort, Karen Fisher-Vanden, David Popp, and Nidhi Santen**, “Should We Give Up after Solyndra? Optimal Technology R&D Portfolios under Uncertainty,” *Journal of the Association of Environmental and Resource Economists*, September 2017, 4 (S1), S123–S151.

- Weiss, Martin, Andreas Zerfass, and Eckard Helmers**, “Fully Electric and Plug-in Hybrid Cars - An Analysis of Learning Rates, User Costs, and Costs for Mitigating CO<sub>2</sub> and Air Pollutant Emissions,” *Journal of Cleaner Production*, March 2019, *212*, 1478–1489.
- Whittle, Peter**, “Restless Bandits: Activity Allocation in a Changing World,” *Journal of Applied Probability*, 1988, *25*, 287–298.
- Wiesenthal, Tobias, Arnaud Mercier, Burkhard Schade, Hrvoje Petric, and Paul Dowling**, “A Model-Based Assessment of the Impact of Revitalised R\&D Investments on the European Power Sector,” *Renewable and Sustainable Energy Reviews*, January 2012, *16* (1), 105–112.
- Wilson, Charlie**, “Up-Scaling, Formative Phases, and Learning in the Historical Diffusion of Energy Technologies,” *Energy Policy*, June 2012, *50*, 81–94.
- , **Arnulf Grübler, Nico Bauer, Volker Krey, and Keywan Riahi**, “Future Capacity Growth of Energy Technologies: Are Scenarios Consistent with Historical Evidence?,” *Climatic Change*, May 2013, *118* (2), 381–395.
- Winkelmann, Ricarda, Jonathan F. Donges, E. Keith Smith, Manjana Milkoreit, Christina Eder, Jobst Heitzig, Alexia Katsanidou, Marc Wiedermann, Nico Wunderling, and Timothy M. Lenton**, “Social Tipping Processes towards Climate Action: A Conceptual Framework,” *Ecological Economics*, February 2022, *192*, 107242.
- Wiser, Ryan, Mark Bolinger, and Eric Lantz**, “Assessing Wind Power Operating Costs in the United States: Results from a Survey of Wind Industry Experts,” *Renewable Energy Focus*, September 2019, *30*, 46–57.
- , – , and **Joachim Seel**, “Benchmarking Utility-Scale PV Operational Expenses and Project Lifetimes: Results from a Survey of U.S. Solar Industry Professionals,” Technical Report None, 1631678, ark:/13030/qt2pd8608q June 2020.

**WMO and IPCC, eds, *Climate Change The IPCC Response Strategies* 1990.**

**— and —, eds, *Climate Change The IPCC Scientific Assessment*, Cambridge University Press, 1990.**

**Wright, Theodore Paul, “Factors Affecting the Cost of Airplanes,” *Journal of the Aeronautical Sciences*, February 1936, 3 (4).**

**Yue, Xiufeng, Steve Pye, Joseph DeCarolis, Francis G.N. Li, Fionn Rogan, and Brian Gallachóir, “A Review of Approaches to Uncertainty Assessment in Energy System Optimization Models,” *Energy Strategy Reviews*, August 2018, 21, 204–217.**

# Chapter III

## The need for better statistical testing in data-driven energy technology modeling

C. LENNART BAUMGÄRTNER, RUPERT WAY,  
MATTHEW C. IVES, J. DOYNE FARMER

Technology modeling is a vital part of developing and understanding energy system scenarios and policy, but it is challenging due to data limitations, deep uncertainty, and the complex social and technological dynamics involved in the evolution of energy systems. These difficulties are often compounded by unsound technology forecasting practice, including over-fitting, data selection bias, and ad-hoc assumptions, leading to unreliable conclusions. We flag several cases where this has been problematic and analyze in detail a recent model for predicting the pace of solar PV and wind energy deployment. We discuss the general takeaways and provide suggestions on how statistical testing can be conducted to avoid such problems in the future and to quantify the reliability of forecasts.

### III.1 Introduction

Decarbonizing our energy system requires dramatically changing the technologies we use to produce, store, and utilize energy. To do this, we need to cope with uncertainty about future technology costs and any potential limits to their deployment. In particular, in order to make sound investments and effective public policy, we require reliable technology forecasting models whose accuracy can be quantified using standard statistical methods ([Grubb et al., 2022](#)). Unfortunately, such models are all too often built without proper statistical testing and with insufficient effort made to understand the reliability of their

---

**Acknowledgments:** For helpful comments and feedback, we thank François Lafond. We also thank Aleh Cherp and coauthors for a fruitful discussion of their earlier work.

conclusions.

One symptom of this problem is the tendency to refer to the outputs of technological change models as “projections”. This word is usually invoked to signal a lack of confidence in the predictive power of the model. Nonetheless, there are many circumstances where projections are used to inform wider energy system scenarios or policy decisions, meaning that they are de facto conditional forecasts about cause and effect, even if they are not labeled as such. Without applying the same statistical testing and validation procedures that are routine in disciplines such as finance, economics, epidemiology, and meteorology to such forecasts, it is impossible to know how reliable we should expect forecasts to be.

Here, we focus on data-driven, statistically-based technological change models that generate forecasts using historical deployment data, historical cost data, or some combination of the two. We will not address equilibrium models (such as most integrated assessment models) or bottom-up engineering models (Krey, 2014). Since data-driven technological change models are often used both to calibrate large energy system models and to inform the wider discourse on scenario feasibility, proper statistical testing and validation are important.

We begin by focusing on the deployment of solar photovoltaic (PV) and wind, which will be the main topic in this paper. Solar PV and wind play an important role in most national and global future energy system scenarios, such as those presented in the IPCC’s most recent assessment report, AR6 (IPCC, 2022). What insights can historical data provide about the likelihood of meeting the deployment levels in different scenarios and the cost of achieving such scenarios? As shown in Figure 1, over the last 35 years, wind electricity production has grown at an average rate of 21% per year, and over the last 45 years, solar PV has grown at an average rate of 32% per year. These rapid growth rates will necessarily diminish through time and eventually level off, but how and when will this happen? Their future deployment rates will most likely be a key determinant of whether or not we can meet the Paris Agreement (UNFCCC, 2015), and predictions about this are essential for planning to meet these goals.

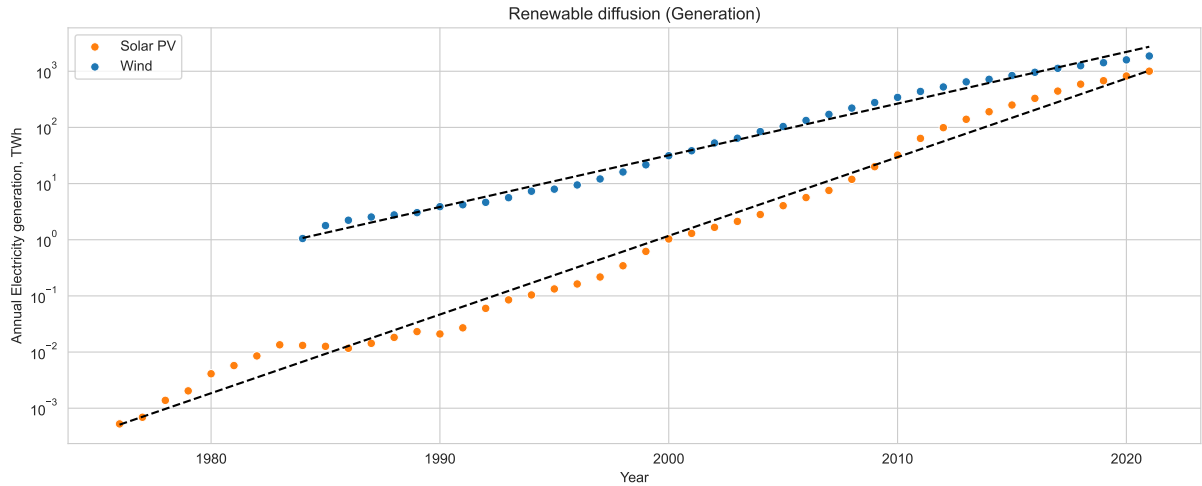


Figure III.1: **Global electricity generation from wind (onshore and offshore) and solar PV.** The historical annual generation of electricity from wind and solar PV shows approximately exponential growth (Way et al., 2022; IRENA, 2022).

Technology deployment trends are frequently characterized by a standard family of models called S-curves. In the early stage of development, absolute deployment levels are small, but growth is rapid. As a technology matures, its growth slows and flattens until it reaches market saturation. This inevitably happens due to technology competition, consumer preferences, and other constraints. For solar PV and wind, slower growth could also occur for socio-political reasons, such as land-use restrictions, public opposition, incumbent interests, insufficient investment, inadequate policy and institutional support, or a lack of cheap energy storage technologies to manage intermittency. However, while it is easy to imagine technology-specific barriers such as these, it is often hard to predict our ability to find solutions, and there is no reliable evidence regarding the timing or size of potential impacts (for example, until recently, grid-scale electricity storage was considered a severe constraint on variable renewables, but this is changing due to plummeting battery costs).

There are many functional forms that can be used to represent S-curves, the simplest example of which is the logistic function. It is written

$$x_t = S(t) = \frac{L}{1 + e^{-k(t-t_0)}}, \quad (\text{III.1})$$

where  $t$  is time in years,  $x_t$  is the number of units of the technology deployed at time  $t$ ,  $L$  is the asymptotic diffusion level,  $k$  is the initial annual exponential growth rate, and  $t_0$  is a location parameter centering the S-curve in time. In the limit  $t \ll t_0$ ,  $S(t) \approx Le^{k(t-t_0)}$ , and in the limit  $t \gg t_0$ ,  $S(t) \approx L$ . The growth rate  $dS/dt$  initially increases exponentially, then begins to noticeably slow down after reaching its maximum at the inflection point  $t = t_0$ , and decreases dramatically as  $x_t$  approaches its asymptotic value  $L$ . This functional form has been shown to provide a good representation of observed data in a variety of technologies, from large infrastructure projects (roads, canals, etc.) to consumer goods (mobile phones, internet users, etc.) (Grübler and Nakićenović, 1991). There are also theoretical reasons for technology adoption following logistic S-curves, based on the study of heterogeneous firms adopting innovations. This is similar to bacteria growth in a Petri dish, where a firm's proximity to a new technology (in terms of other firms having adopted it) increases the likelihood of the firm adopting the innovation (Fisher and Pry, 1971). It is therefore reasonable to expect energy technologies to follow a similar pattern. However, while S-curves often provide a good fit to observed data after a technology has already reached global saturation, real-world data never follows an S-curve exactly, and as yet S-curve based methods have not been tested to assess their performance as forecasting tools for technologies in general.

Modeling the future diffusion of novel low-carbon technologies is challenging, given the complex social and technological dynamics associated with the energy system (Farmer et al., 2015; Geels et al., 2017). The pace of change means that the clean technologies we are most interested in have limited historical records, and their socio-economic environment is rapidly evolving. This poses an issue for all forward-looking models, especially for data-driven approaches relying on historical observations for model calibration and validation. However, this does not mean that S-curve models are useless in this context; rather, it means that the reliability of their predictions needs to be assessed carefully so we know how useful they are. As discussed later, for technologies like solar PV and wind that are still in the early stages of their development, due to data limitations, it is necessary to

test any new prediction methodologies on many different technologies, particularly those that have already reached maturity. This provides ground truth for testing predictions in situations where the correct answers are known.

Data limitations mean that it is very easy to *overfit* historical data, i.e., to choose models that match the noisy features of the data rather than the underlying signal. This typically occurs when the data are noisy, and the number of data points is small compared to the number of parameters. For S-curves, avoiding overfitting is particularly challenging due to their highly non-linear nature. When forecasting incomplete S-curves, the model is frequently overfit, where temporary slowdown in adoption is misinterpreted as the long-term asymptote (Petropoulos et al., 2022). Other important problems are *selection bias*, where the selective exclusion of data distorts results, and the use of unfounded *ad-hoc assumptions* that are chosen for convenience rather than being based on evidence (see Petropoulos et al. (2022) for a comprehensive overview of data-driven predictive models and their challenges). Unfortunately, these mistakes are often made in statistical technological change models, resulting in poor predictions.

In this paper, we primarily focus on the problem of forecasting technology deployment and mention several examples that illustrate such mistakes. We also provide an example from cost forecasting to demonstrate that such problems are not unique to the prediction of deployment. We begin by discussing a range of examples where predictions were unreliable in hindsight, then dive deeply into one specific example as a case study. We conclude by discussing the issues of overfitting, selection bias, and poor statistical practice more generally, as well as providing some suggestions about making more reliable models and assessing the quality of their predictions.

## III.2 Examples of overfitting and selection bias

The difficulty of forecasting technology deployment directly using the logistic function is illustrated in Figure III.2, which shows point forecasts for the deployment of solar PV (left panels) and wind (right panels), in terms of annual electricity generation (top panels) and

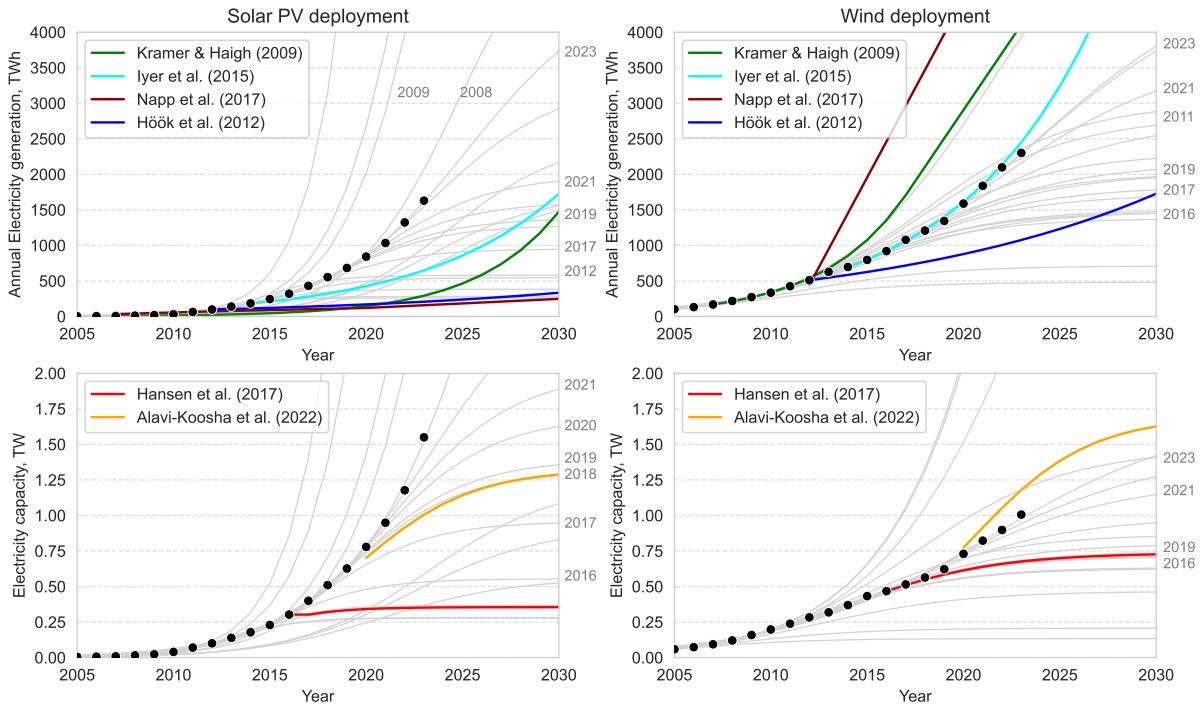


Figure III.2: **Global deployment for solar and wind (onshore plus offshore) vs. past deployment projections.** Annual electricity generation (top row) and generating capacity (bottom row) for solar (left) and wind (right). **(dots)** Actual annual electricity generation and capacity. Data prior to 2005 is not shown here for better readability. **(gray lines)** S-curve fits using the logistic function described in the main text, for  $t_{max}$  between 2005 and 2019. We see that S-curves are highly sensitive to  $t_{max}$ , as indicated by the gray labels. Recent fits align more closely with the data; however, it is unknown how well they will align with future deployment. **(colorful lines)** Predictions for future solar and wind deployment in the literature. Similar to the S-curve fits, future deployment predictions vary widely in the literature. For solar, all predictions under-estimate future deployment. Data until 2019 based on [IEA \(2023\)](#), until 2023 based on [Ember \(2024\)](#).

generating capacity (bottom panels). We pretend to be at a time  $t_{max}$  in the past, fitting logistic S-curves on all data prior to  $t_{max}$ . The gray lines in Figure III.2 show logistic S-curve fits for different values of  $t_{max}$ . As the figure illustrates, the fits are highly sensitive to  $t_{max}$ . The fits for solar vary from predicted asymptotic capacities as low as a quarter of a terawatt (a factor of six lower than today’s capacity) to extremely high values that are too large to be seen in the plot. In later years, the fits are less sensitive to  $t_{max}$  but remain unstable. Furthermore, the fits in recent years tend to underestimate future diffusion, as is evident by the fact that most of them are to the right of the real data.

The inconsistent S-curve fits in Figure III.2 provide a good example of overfitting. Although the logistic function may seem simple, it is a nonlinear function with three free parameters, and even small amounts of noise can make fits unstable until diffusion is

sufficiently developed, which is not true yet for either solar or wind. For example, we observe a slowdown in wind deployment 2012 to 2019. As a result, S-curve fits during that period substantially underestimate wind deployment after 2020. Note that before 2012, wind deployment was both over- and under-estimated as evident from the additional gray lines in Figure III.2.

Despite this, S-curve fits have been used multiple times in the literature to forecast future renewables growth. For example, Hansen et al. (2017) (Figure III.2, red line) fitted S-curves to data up until 2015 and concluded that “the logistic model implies that the total installed [wind and solar PV] capacity saturates at around 1.8 TW in 2030.” Even though these claims have since been refuted methodologically (Rypdal, 2018), and the fact that by 2023 we had already installed 2.4 TW of solar PV and wind (IRENA, 2024), such statements are still persistent in the literature (e.g., Madsen and Hansen (2019)).

Another example of likely overfitting is given in Alavi-Koosha et al. (2022), who forecasted future renewable capacities using a weighted average of four different kinds of S-curves (Figure III.2, orange line). While weighting different models can lead to better forecasts, it is not a panacea – if all the forecasts are bad, or if there are systematic biases, the resulting forecasts may still be bad. As shown in Figure III.2, the forecasts have been significantly too low for solar PV and too high for wind, indicating that the underlying model is likely overfit.

Perhaps more common than overfitting is the issue of selection bias. We have found a number of examples in the literature that have analyzed historical growth rates in (energy) technologies to imply limits for future renewable growth rates.

In 2009, Kramer and Haigh (2009) (Figure III.2, green lines) used historical data from past energy technologies to argue that annual growth in clean energy technologies is likely to be at most 26% during their early deployment phase, until they achieve “materiality” (specified as “around 1%” of global primary energy production), after which annual growth will slow dramatically and proceed linearly. However, if we assume 26% annual growth from 2009, materiality will not happen before 2030 for solar PV or before 2017 for wind.

As shown in Figure III.2, these limits provide a reasonable upper bound for wind but have already been surpassed by solar PV.

Höök et al. (2012) also studied growth rates in previous energy technologies (nuclear, liquid natural gas, biofuels, wind) and concluded that wind and solar PV electricity generation are unlikely to grow faster than 7% per year (Figure III.2, blue line). This is based on the fact that fossil fuels have never grown faster than this rate over multiple decades. This comparison was repeated in a similar fashion by Smil (2016), even though prior renewable growth had already been significantly larger than 7%.

Iyer et al. (2015) built on the work of Kramer and Haigh and Höök et al., to examine the historic deployment rates of around 40 technologies, including non-energy technologies, such as washing detergent and black-and-white televisions. They concluded based on their wider subset of technologies that annual growth rates above 15% were unlikely for renewables, and applied this limit to an IAM (GCAM) (Figure III.2, cyan lines). While adequate for wind electricity, this also under-estimated future growth in solar PV.

Napp et al. (2017) performed a similar analysis, limiting renewable growth to 20% by the same arguments. They applied these limits to another IAM (TIAM-Grantham) (Figure III.2, brown lines). Again, this upper bound might be reasonable for wind but substantially underestimates growth in solar PV.

In all of these examples, researchers investigated a sample of historic technology growth trajectories and then inferred limits on future renewable growth. While mostly adequate for wind, all of these limits were too low for solar PV, as evident from Figure III.2. It appears that these studies were all subject to selection bias; their conclusions are biased because too few technologies were considered in their analyses. For example Iyer et al. (2015) include only 4 out of 38 technologies with annual growth rates above 15% (cars, locomotives, wind energy, washing detergent) and no technology above 24% annual growth rate. This means they exclude other technologies, such as mobile phones and internet users, both of which grew at rates up to 30% annually. Whilst it is unclear which of these technologies can serve as a good direct comparator to wind and solar, excluding them

from more holistic analyses a priori can lead to selection bias.

Issues with overfitting and data selection bias also occur when modeling the cost of technologies. For example, [Hsieh et al. \(2019\)](#) applied a two-stage learning model to lithium-ion battery prices, in which raw material prices were treated separately from material synthesis and battery pack manufacturing costs. They modeled raw material prices of lithium, cobalt, and nickel by extrapolating from a very limited selection of the historical price record, and predicted that in 2023 nickel-manganese-cobalt (NMC) battery pack prices would be in the range 170-201 USD(2023)/kWh. However, cobalt prices have dropped significantly since their publication, and in 2023 the average NMC battery pack price was at an all-time low of 146 USD(2023)/kWh ([Catsaros, 2023](#); [Gül et al., 2024](#)). The average price for battery packs used in electric vehicles was lower, at 128 USD(2023)/kWh ([Catsaros, 2023](#)), due to a shift towards lower cost lithium-iron-phosphate batteries, and all pack prices have continued to decline in 2024 ([McKerracher, 2024](#)). This indicates that their forecasts were likely overfit. Subsequently, [Penisa et al. \(2020\)](#) tested a similar learning model for predicting EV battery costs and rejected a model that includes lithium and cobalt metal prices on the basis that it was overfit [Nagy et al. \(2013\)](#). showed that multi-factor learning models tend to overfit data because they have too many parameters and, therefore, provide worse forecasts of future costs than simple alternatives. The approach developed by ([Hsieh et al., 2019](#)) is nevertheless being used to forecast battery prices in the U.S. decades into the future ([EIA, 2023](#)).

### **III.3 Testing an energy technology diffusion model from the recent literature**

To highlight how overfitting and selection bias can affect the reliability of a statistical technological change model, we consider in detail a recent example from the literature. [Cherp et al. \(2021\)](#) developed a model of deployment growth in solar PV and wind to assess the feasibility of their rapid deployment in energy scenarios. Based on historical

growth in these two technologies, they conclude that “some 1.5° and 2°C pathways pose serious feasibility concerns” and “replicating or exceeding the fastest national [historic] growth may be challenging.”

These conclusions have attracted significant attention in the academic community and wider media (Nature Energy, 2024), and their methodology has been replicated to comment on feasible national trajectories for solar and wind (Jakhmola et al., 2023), nuclear (Vinichenko et al., 2021), and carbon direct removal (CDR) (Nemet et al., 2023). In the following, we discuss their method and explain why it suffers from problems of overfitting and selection bias.

The method of Cherp et al. (2021) aims to take advantage of the fact that some countries already have substantial deployment of solar and wind energy while others have little to none. They fit S-curves to historical diffusion data in 60 countries to estimate what they call the *maximum historic growth rate*  $G$  for each technology in each country.  $G$  is calculated as the maximum growth rate of the S-curve fitted to the annual electricity production time series of a given country, measured relative to that country’s total electricity supply. ( $G$  should not be confused with the annual exponential growth rates quoted earlier.)

The maximum growth rate  $G$  of the logistic S-curve occurs at the inflection point  $t_0$ , where  $G = kL/4$  (normalized by the total electricity supply at  $t_0$ ). Based on S-curve fits, Cherp et al. (2021) identify what they believe are countries that have already passed their period of maximum growth. This is done by comparing the current deployment level  $x_t$  to the fitted (logistic) asymptotic market size  $L$ . For a given country at time  $t$ , they deduce that it has already passed its period of maximum growth if the maturity  $m = x_t/L > 0.5$ . (The value 0.5 is a natural choice as it corresponds to the diffusion level relative to the asymptote  $L$  at the inflection point  $t_0$ .) Cherp et al. (2021) then compare their estimates of  $G$  for these “high maturity” countries to growth rates required for solar and wind in rapid decarbonization scenarios. (See Appendix A.1.1 for a detailed description of the method.) Because their estimates of the maximum growth rates are lower than those

required in 1.5° or 2°C scenarios, they conclude that reaching our climate goals will be “challenging”.

Overfitting occurs when data are well-matched by a model in-sample but poorly matched out-of-sample. Here, *in-sample* refers to data used to calibrate a model, meaning to estimate its parameters or perform model selection, and *out-of-sample* refers to data that was not used for calibration. Since [Cherp et al. \(2021\)](#) only analyzed data up to and including 2019, we can test for overfitting by using more recent data. [Figure III.3](#) compares actual vs. predicted maximum growth rates for individual countries in and out-of-sample. Rather than relying on an S-curve fit, we directly estimate the maximum growth rate observed at any point in time by taking the maximum average growth rate over all possible historical three-year intervals. (Three-year intervals are chosen to reduce noise.) We then compare our direct measurements to the maximum historical growth rates estimated by [Cherp et al. \(2021\)](#). When doing this, we only consider countries that [Cherp et al. \(2021\)](#) believe to have already reached or passed their point of maximum growth. The out-of-sample comparison is based on data through 2022 for solar PV and 2021 for onshore wind.

[Figure III.3](#) plots the growth rates observed in each year against those predicted by [Cherp et al. \(2021\)](#). The in-sample growth rates in 2019, the year of their study, are clustered along the identity line, as expected. This is not the case, however, for later out-of-sample growth rates. All but three of the out-of-sample data points (shown as blue dots) are above the identity line, and many are substantially above it. While [Cherp et al. \(2021\)](#) predict that growth rates above 2% for solar and 3.5% for wind to be rare, the observed maximum growth rates are as high as 3.3% for solar and 4.4% for wind. This is a clear demonstration of overfitting.

As mentioned in the introduction, in situations where the data are severely limited, such as for solar and wind diffusion, it is essential to perform tests on other technologies. This is also required to assess whether model dynamics are likely to be stationary and, if not, what the effects might be. The S-curve model is *technology agnostic*, meaning that




Figure III.3: Maximum historic growth rates  $G$  for solar photovoltaics (left) and onshore wind energy (right) as calculated by [Cherp et al. \(2021\)](#) versus maximum 3-year moving average growth rates until 2015 (red dots), 2019 (orange triangles), and 2022/2021 (blue dots). The observed growth rates for each country as a percentage of its annual electricity supply are calculated as the maximum three-year moving average of the annual growth rate ([Cherp et al., 2021, Ext. Fig. 7b](#)) and plotted against the maximum historic growth rates  $G$  as calculated through S-curve fits performed by [Cherp et al. \(2021\)](#). For solar PV, the red dots, orange triangles, and blue dots consider the full national time series until 2015, 2019, and 2022, respectively. For onshore wind, the red dots, orange triangles, and blue dots consider the time series until 2015, 2019, and 2021, respectively, due to data availability. For countries with significant electricity generation by offshore wind, we used data until 2020 since the 2021 data were not yet available. (The sample size has dropped to 17 countries for solar PV and 27 countries for onshore wind because a majority of countries are excluded by the filtering method of [Cherp et al. \(2021\)](#).) The 2019 data that were used to produce the fitted values of  $G$  are consistent with their predictions, but later out-of-sample data substantially exceed the predicted maximum rates. Similarly, earlier data fall below the predicted maximum rates. Data for solar is from [Ember \(2024\)](#), and data for wind is from [IEA \(2023\)](#); [IRENA \(2022\)](#).

all technologies follow some form of an S-curve, and it can be applied to any technology ([Grübler and Nakićenović, 1991](#)). As with the logistic S-curve model, there is nothing about [Cherp et al. \(2021\)](#)'s method that is specific to solar or wind energy. This can be used to vastly enlarge the quantity of data available for testing. Indeed, pooled technology datasets have already been extensively used in developing cost forecasting models ([Nagy et al., 2013](#); [Farmer and Lafond, 2016](#); [Lafond et al., 2018](#)). We take advantage of the fact that the method of [Cherp et al. \(2021\)](#) is technology agnostic to test it on the global level and get more insight into its reliability. For this purpose, we choose mature technologies that are already close to asymptotic deployment and so provide ground truth. Thus, unlike solar and wind energy, which are still at a stage where their asymptotic levels of deployment are unknown, we can apply their method during different stages of a technology's development

and then compare the results to what really happened. In the main text, we present the results of a back-test for mobile phone cellular subscriptions (ITU, 2022), and in the Supplementary Materials A.1.3, we present the same analysis for annual airline passenger miles and internet users, which all provide similar results.

A visual examination of the mobile phone data in Figure A.2 makes it clear that by 2020, the technology diffusion process for mobile cellular subscriptions was well past its inflection point, with the S-curve stagnating at the global and most national levels. Fitting a logistic S-curve to the global data yields a maximum growth rate  $G$  of roughly 9% at the inflection point  $t_0 \approx 2008$  and an asymptote  $L \approx 123\%$  of the total population. The estimated initial growth rate per year for mobile phones is  $k \approx 30\%$ , which is comparable to the annual growth rates for solar and wind energy. In the same way that Cherp et al. (2021) normalize solar and wind deployment to total electricity supply, we normalize the number of phone subscriptions to the national (global) population. We use their code and data selection procedure to calculate the distribution of maximum national historical growth rates  $G$  for all possible maximum years  $t_{max}$ .

The change in the distribution over time is shown in Figure III.4. The dots are the national maximum growth rates  $G$  for countries with  $m > 0.5$ . The solid line shows the median of the resulting distribution, and dotted lines indicate interquartile ranges (25%, 75%). The figure makes it clear that the distribution of  $G$  strongly depends on the choice of  $t_{max}$ . The change in the median is indicative of the change in the distribution as a whole. The distribution briefly appears (roughly) stationary with a median of about 1% in the early phase running from  $t_{max} = 1989$  to  $t_{max} = 1993$ . After  $t_{max} \approx 2005$ , the distribution is roughly stationary again, but the median is more than a factor of 10 higher than it is in the early stage. The change happens abruptly: When  $t_{max} = 1995$ , the median of  $G$  is below 2%; when  $t_{max} = 1998$ , it is about 4%, and when  $t_{max} = 2000$ , it is about 10%, which is still only around 75% of the final value of 13.3%. (The fact that this is higher than the estimated growth rate of 9% above is due to variations in the S-curve fits with different final years). Prior to reaching global maturity, the three technologies that we test

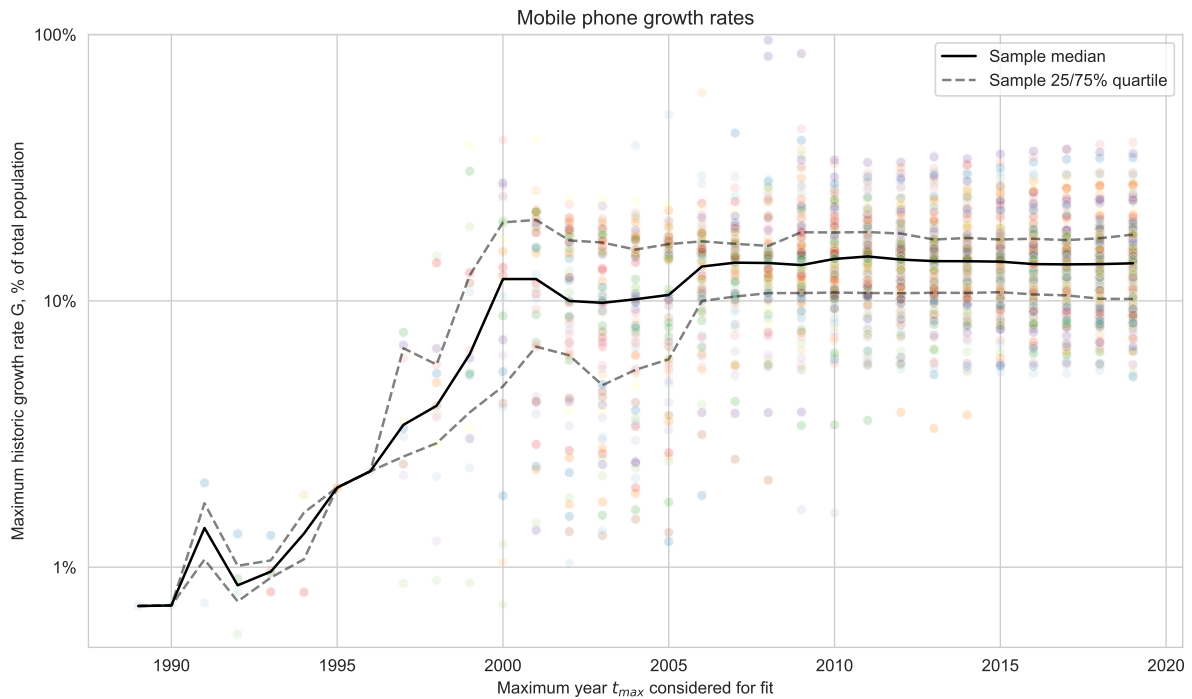


Figure III.4: **Distribution of national maximum historic growth rates ( $G$ ) for mobile phones vs. the maximum year considered for the S-curve fit ( $t_{max}$ ).** National maximum historic growth rates included in the sample according to the filtering applied by Cherp et al. (2021) are shown as dots, each color corresponding to one country. The solid line shows the median maximum historical growth rate in the sample. Dashed lines represent the interquartile range of the distribution. This illustrates the non-stationarity of the distribution of  $G$ , and that early forecasts based on the Cherp et al. (2021) method yield results that are too low by more than an order of magnitude.

all exhibit non-stationary behavior in the distribution of  $G$ , only stabilizing once they are close to maturity. This indicates that, once the dust settles, it is likely that we will see similar results for wind and solar.

Both of the tests for overfitting we have done so far relied on gathering new data. Based on the principles of backtesting, there are also opportunities to test for overfitting without adding new data. In Figure III.3, the non-stationarity of the results is also visible by comparing predictions in 2019 to earlier measurements of the maximum growth rates. The red dots in Figure III.3 correspond to the maximum growth rates measured in 2015. These are mostly below the identity line, demonstrating that the maximum growth rates still tend to grow with time, as one would expect for countries that have still not yet reached their true maximum growth rate.

Sensitivity testing could also have been used to make the underlying problem clear.

This is demonstrated in Figure III.5, where we compare fits for onshore wind power deployment in Austria, China, and Finland. We fit S-curves to the data from the first data point up to a maximum time  $t_{max}$ , which we vary. As shown in Figures III.5a-c, for these countries, the resulting fits for  $t_{max} = 2012$  are wildly different from those for  $t_{max} = 2019$ . As seen in Figures III.5d-e, the fitted S-curve parameters also vary greatly. In Figure III.5e, the estimates of the asymptotic deployment rate  $L$  reach levels as high as 3,500% of the national electricity supply. Similarly, the exponential growth rate  $k$  in Figure III.5d shows large fluctuations between 120% and 20%. As a result, the estimated maximum growth rates ( $kL/4$ ) seen in Figure III.5g reach levels as high as 400% per year.

Cherp et al. (2021) make an effort to manage these instabilities by only considering countries that the S-curve fit suggests have already passed their point of maximum growth. This is done based on the requirement that  $m = x_{t_{max}}/L > 0.5$ . However, as we see in Figures III.5a-c, this often results in countries going in and out of the selected sample of “mature” countries. Figure III.5f shows the maturity  $m$  for Austria, China, and Finland for different years of evaluation  $t_{max}$ , with the cutoff  $m = 0.5$  as a reference. For their method to work as intended, one expects that once a country is deemed “mature”, it should generally remain in the sample, but this is not the case. Finland, for example, is classified as mature before 2010, then immature from 2011–2017, then mature from 2018 onward. Austria behaves similarly, leaving the sample in 2015 after many years of being included. China goes back and forth between included and excluded between 2007 and 2014.

In Figure A.1 of the Supplementary Materials, we show that we have not cherry-picked these examples. The estimated year  $t_0$  – where national growth should reach its maximum – increases systematically with time. On average, when  $t_{max}$  increases by a year, the mean value of  $t_0$  increases by 0.6 years for wind and 0.3 years for solar. Across the length of the samples, the mean value of  $t_0$  increased by 8.4 years for wind and 2.1 years for solar. This demonstrates that the answers yielded by their procedure are strongly non-stationary.

These problems stem from a combination of overfitting and selection bias. The overfitting occurs because there is not enough data and too much noise to fit a nonlinear

function with three parameters. This means that the estimation of the parameters of the S-curve is not statistically significant, which in turn means that the determination of maturity based on the criterion  $m = x_{t_{max}}/L > 0.5$ , is not statistically significant either. The lack of statistical significance means that the selection of mature vs. immature countries is inconsistent. The fluctuations that indicate slow growth pass their test, while those that indicate rapid growth fail it, introducing selection bias. Only later, when the S-curves have reached maturity, do the fits become statistically significant, meaning that almost all countries pass the maturity test, and the sampling bias is removed.

### III.4 Key Takeaways: Making reliable statistical inferences

We now summarize some of the key takeaways and suggest ways to avoid such problems in the future to produce more reliable technological change models.

The first takeaway is that *statistical testing* is essential. In all of the examples discussed here, statistical testing might have provided criteria for determining how reliable the models were. In the example of [Cherp et al. \(2021\)](#), assessing the statistical significance for their fits when judging whether a country has reached maturity could have made their conclusions more robust. However, there are two big challenges in doing so: The data for national diffusion are noisy, and the time series are short. There is a risk that under a meaningful test of statistical significance in measuring maturity, very few countries would have maturities with enough statistical significance to be useful at the present time.

This brings up another important lesson: *Some questions cannot be answered*. There are many circumstances where there is insufficient data to answer questions quantitatively. In instances where statistical significance can not be estimated, any answer needs to be treated with extreme caution. It is much better to say “we cannot reach a conclusion” rather than reaching a misleading conclusion.

An important caveat to statistical testing, though, is that it requires great care. While many canned routines for fitting functions to data provide confidence intervals, these typically make strong assumptions, for example, the assumption that the data are

Independent and Identically Distributed (IID), which is often not appropriate for real data. For example, technology diffusion trajectories are usually strongly autocorrelated, meaning that the residuals from the S-curve fit are not independent, and the fluctuations tend to be very persistent. Moreover, S-curve growth is inherently non-stationary – the behavior during the early stages of a technology’s evolution is intrinsically different than the behavior during later stages. In part because of this, technology trajectories are heteroscedastic, meaning that the amplitude of the residuals varies with time, and consequently, the residuals are not identically distributed. These effects reduce the number of independent degrees of freedom and thus reduce statistical significance. As a result, methods assuming IID noise will often give over-optimistic answers that impart a misleading sense of confidence. Getting around this problem using in-sample tests is difficult and typically requires advanced numerical methods. (For an example of how statistical testing can be done when predicting technology costs, see [Farmer and Lafond \(2016\)](#).)

One of the most reliable ways to cope with this problem is to test conclusions out-of-sample. This is typically done by dividing the data into two or more statistically independent samples, performing model selection and parameter estimation on one (training) sample, and testing the resulting model on the other (testing) sample. This procedure has the huge advantage that it is not necessary to characterize exactly how the data deviates from being IID, which is usually difficult. It directly tests for overfitting, and it automatically takes deviations from IID data into account. Furthermore, it provides a straightforward way to measure how good the predictions are, which is valuable by itself: Predictions without error estimates attached are less useful.

Though testing out-of-sample is the most reliable way to test for statistical significance, it still has pitfalls. While developing a model, researchers usually test many different models and select the one with the best performance. When tests are made repeatedly, this procedure becomes unreliable because the model’s performance becomes positively biased, i.e., the model’s score is inflated. The only test that is truly out-of-sample is the first one: As soon as a second test is made, the data are no longer out-of-sample – they have become

part of the training set. To cope with this problem, it is good practice to divide the data into three samples for training, testing, and final validation. The training set is used to estimate the parameters of a model, and the test set is used for model selection. The final validation set is used only after a model has been selected and then can only be used once. If the final validation fails, repeated model selection comes with the danger of overfitting. Because of this, it is helpful for researchers to be very clear about their model selection protocol and ideally record the number of times that out-of-sample tests were carried out.

Data selection bias can emerge when data is selectively excluded, either from the model calibration process or in post-processing, i.e., after the model parameters have been fit. Reasons for excluding data are often justified, such as coping with obvious data errors or excluding outliers to make parameter estimation more robust. In technological change studies, researchers may also want to exclude very old data points (or give them lower weight) because of heteroscedastic noise or structural breaks. Data selection can also occur implicitly when a dataset does not include certain data, for example, if data from developing countries is missing from a global study. The worst problems with data selection occur when selection causes systematic effects so that the data that are excluded have different properties than those that are not excluded. Such problems can be subtle, and avoiding them requires careful consideration of any procedures involving data selection. In the case of S-curves, it is unclear if and how historical data should be excluded. Multiple studies, including [Cherp et al. \(2021\)](#) have suggested excluding early data points due to heteroscedastic noise. This may be justified but requires a more in-depth analysis of the noise model to provide reliable insights.

Finally, there are two fundamental reasons why models do not replicate reality exactly: the model bias and variance. A typical modeling exercise uses a functional form or model family (like the logistic S-curve) with free parameters (like  $k$ ,  $t_0$ , and  $L$ ) to fit the data. (For this discussion, we use the term “model” to refer to a specific model with fixed parameters and “model family” to refer to a set of possible models determined by the free parameters, like the logistic S-curve). *Bias* refers to the situation where no parameters of the model

family exist that can match the real data. For example, the exponential model family will most likely provide biased forecasts for future technological adoption since it does not include an asymptote. A common way to cope with bias is to introduce more parameters into the model family. However, this usually makes the errors in estimating the parameters worse, a problem that is called *variance*. For example, the logistic S-curve can have lower bias than the exponential model (since it includes an asymptote) but may have higher variance due to its additional parameters. So, while there may exist parameter values that fit the data much better, it is impossible to find them because of larger estimation errors. This causes what is called the bias-variance tradeoff: The best model family is neither too simple nor too complicated; this happens when the errors caused by bias are roughly equal to the errors caused by variance.

Model users unfamiliar with this issue tend to overestimate the problem of bias and underestimate the problem of variance. When datasets are short and noisy, the best models are often much simpler than intuition might suppose. For example, for the problem of predicting technology costs, [Nagy et al. \(2013\)](#) tested various hypotheses against data for 50 technologies; model families with two parameters made the best predictions, while all model families with three free parameters did worse.

To summarize, researchers working on technological change models can reduce overfitting and data selection bias by using statistical testing, ideally out-of-sample. It is important to understand the reliability of predictions – this is usually a much harder problem than making predictions, but predictions of unknown quality are less helpful. Unsound inferences could lead to unsound policy advice, which risks unnecessary delays in the energy transition. Indeed, despite several predictions in the literature that we have highlighted here, there are no statistically sound S-curve forecasting methods that suggest solar and wind energy can not grow fast enough to meet the Paris Agreement targets or that this would be extraordinarily challenging in comparison with past technology growth trends. The procedures that we advocate above are routine in machine learning and for quantitative hedge funds, where survival depends on forecast reliability. The stakes for

climate change are far higher, and the forecasting of climate mitigation deserves at least as much care.

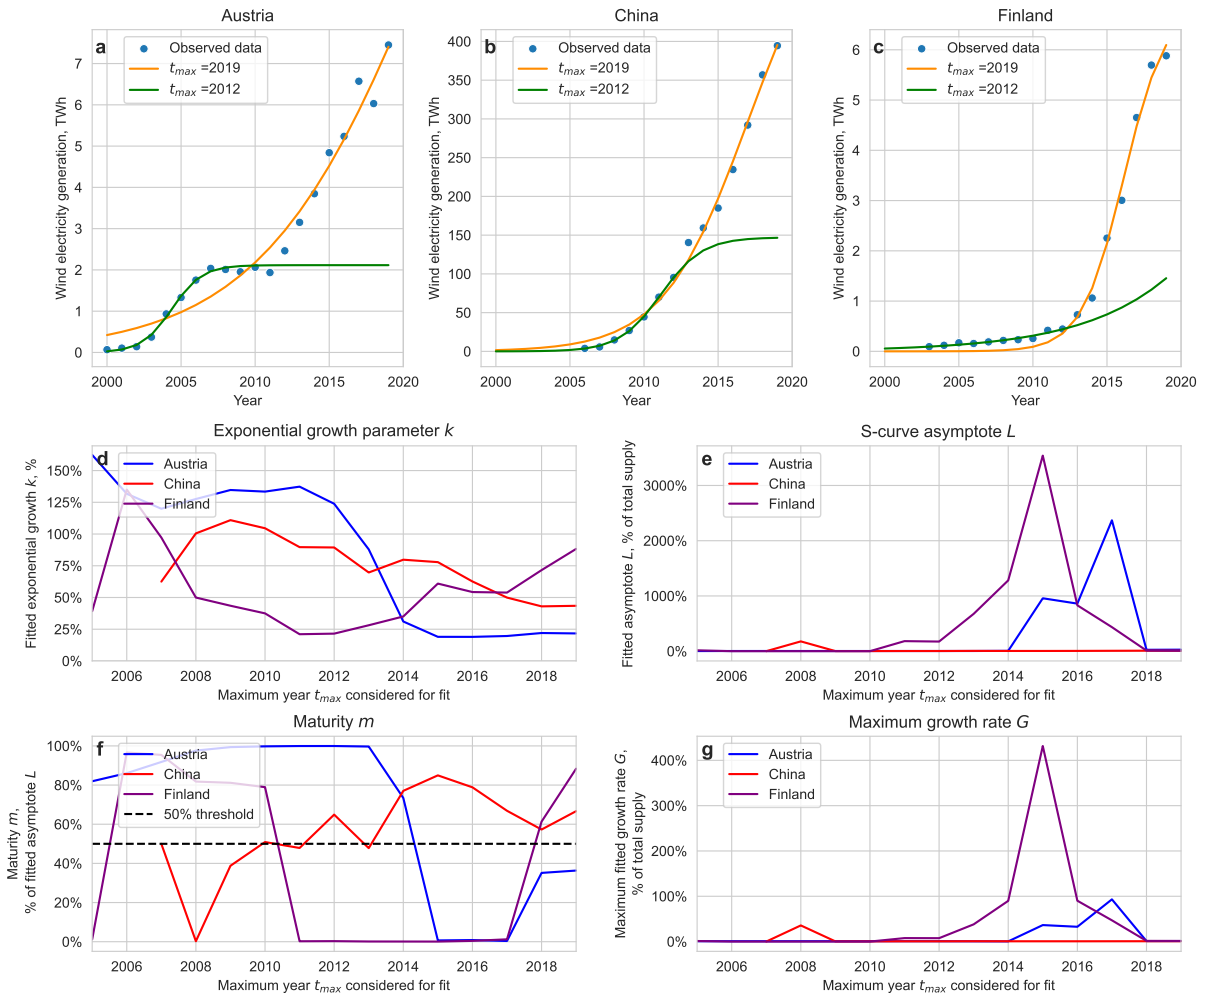


Figure III.5: S-curve fits for electricity generation from onshore wind in Austria, China, and Finland for different maximum years  $t_{max}$ . (a-c) The deployment of onshore wind in (a) Austria, (b) China, and (c) Finland shows intervals of rapid acceleration and deceleration. Consequently, the logistic S-curve fit based on data up until  $t_{max} = 2012$  (green) differs starkly from that including data up until  $t_{max} = 2019$  (orange). Both the exponential growth parameter  $k$  of the resulting S-curve fit (d) and the S-curve asymptote  $L$  (e) vary greatly depending on the length of the time series. As a result, the maturity  $m$  (f) and the maximum growth rate  $G$  (g) are also highly sensitive to the maximum year  $t_{max}$ . Figure A.4 in the Supplementary Materials shows the equivalent figure using the Gompertz model, an alternative S-curve specification.

## References

- Alavi-Koosha, Ahmadreza, Amirhossein Akbari, Mohammadreza Toulabi, and Turaj Amraee**, “Trend Curve- and Machine Learning-Based Renewable Energy Development Forecast,” *2022 12th Smart Grid Conference, SGC 2022*, 2022, (Sgc), 1–5.
- Catsaros, Oktavia**, “Lithium-Ion Battery Pack Prices Hit Record Low of \$139/kWh,” *BloombergNEF*, November 2023.
- Cherp, Aleh, Vadim Vinichenko, Jale Tosun, Joel A Gordon, and Jessica Jewell**, “National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets,” *Nature Energy*, 2021, 6, 742–754.
- EIA**, “Transportation Demand Module Assumptions,” Technical Report March 2023.
- Ember**, “Yearly Electricity Data,” <https://ember-climate.org/data-catalogue/yearly-electricity-data/> 2024.
- Farmer, J. Doyne and François Lafond**, “How Predictable Is Technological Progress?,” *Research Policy*, April 2016, 45 (3), 647–665.
- , **Cameron Hepburn, Penny Mealy, and Alexander Teytelboym**, “A Third Wave in the Economics of Climate Change,” *Climate Research in Gothenburg on*, 2015, 62, 329–357.
- Fisher, Jacob C. and Robert H. Pry**, “A Simple Substitution Model of Technological Change,” *Technological Forecasting and Social Change*, January 1971, 3, 75–88.
- Geels, Frank W, Benjamin K Sovacool, Tim Schwanen, and Steve Sorrell**, “The Socio-Technical Dynamics of Low-Carbon Transitions,” *Joule*, 2017, 1 (3), 463–479.
- Grubb, Michael, Chukwumerije Okereke, Jun Arima, Valentina Bosetti, Ying Chen, James E. Edmonds, Shreekanth Gupta, Alexandre C Köberle, Snorre Kverndokk, Arunima Malik, and Linda Yanti Sulistiwati**, “2022: Introduction

and Framing,” in “IPCC 2022: Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change,” 2022 ed., Cambridge, UK: Cambridge University Press, 2022.

**Grübler, Arnulf and Nebojša Nakićenović**, “Long Waves, Technology Diffusion, and Substitution,” *Review (Fernand Braudel Center)*, 1991, *14* (2), 313–343.

**Gül, Timur, Araceli Fernandez Pales, and Elizabeth Connelly**, “Global EV Outlook 2024 Moving towards increased affordability,” Technical Report, IEA 2024.

**Hansen, Jan P., Patrick A. Narbel, and Dag L. Aksnes**, “Limits to Growth in the Renewable Energy Sector,” *Renewable and Sustainable Energy Reviews*, April 2017, *70*, 769–774.

**Höök, Mikael, Junchen Li, Kersti Johansson, and Simon Snowden**, “Growth Rates of Global Energy Systems and Future Outlooks,” *Natural Resources Research*, March 2012, *21* (1), 23–41.

**Hsieh, I-Yun Lisa, Menghsuan Sam Pan, Yet-Ming Chiang, and William H. Green**, “Learning Only Buys You so Much: Practical Limits on Battery Price Reduction,” *Applied Energy*, April 2019, *239*, 218–224.

**IEA**, “World Energy Balances,” [https://doi.org/10.1787/ data-00512-en](https://doi.org/10.1787/data-00512-en) 2023.

**IPCC**, *Climate Change 2022 - Mitigation of Climate Change - Full Report* number 1, Cambridge University Press, Cambridge, U.K., 2022.

**IRENA**, “Renewable Capacity Statistics 2022,” Technical Report, International Renewable Energy Agency, Abu Dhabi April 2022.

**IRENA**, “Renewable Capacity Statistics 2024,” Technical Report, International Renewable Energy Agency 2024.

ITU, “ITU DataHub,” <https://datahub.itu.int/data/?i=178> 2022.

**Iyer, Gokul, Nathan Hultman, Jiyong Eom, Haewon McJeon, Pralit Patel, and Leon Clarke**, “Diffusion of Low-Carbon Technologies and the Feasibility of Long-Term Climate Targets,” *Technological Forecasting and Social Change*, January 2015, *90*, 103–118.

**Jakhmola, Avi, Jessica Jewell, Vadim Vinichenko, and Aleh Cherp**, “Projecting Feasible Medium-Term Growth of Wind and Solar Power Using National Trajectories and Hindcasting,” July 2023.

**Kramer, Gert Jan and Martin Haigh**, “No Quick Switch to Low-Carbon Energy,” *Nature*, December 2009, *462* (7273), 568–569.

**Krey, Volker**, “Global Energy-Climate Scenarios and Models: A Review,” *Wiley Interdisciplinary Reviews: Energy and Environment*, 2014, *3* (4), 363–383.

**Lafond, François, Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J. Doyne Farmer**, “How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts,” *Technological Forecasting and Social Change*, March 2018, *128*, 104–117.

**Madsen, Dorte Nørgaard and Jan Petter Hansen**, “Outlook of solar energy in Europe based on economic growth characteristics,” *Renewable and Sustainable Energy Reviews*, 2019, *114*, 109306.

**McKerracher, Colin**, “China’s Batteries Are Now Cheap Enough to Power Huge Shifts,” *Bloomberg.com*, July 2024.

**Nagy, Bela, J. Doyne Farmer, Quan M. Bui, and Jessika E. Trancik**, “Statistical Basis for Predicting Technological Progress,” *PLoS ONE*, February 2013, *8* (2), 52669.

**Napp, Tamaryn, Dan Bernie, Rebecca Thomas, Jason Lowe, Adam Hawkes, and Ajay Gambhir**, “Exploring the Feasibility of Low-Carbon Scenarios Using Historical Energy Transitions Analysis,” *Energies*, January 2017, *10* (1), 116.

**Nature Energy**, “Article Metrics - National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets,” <https://www-nature-com.ezproxy-prd.bodleian.ox.ac.uk/articles/s41560-021-00863-0/metrics> 2024.

**Nemet, Gregory F., Jenna Greene, Finn Müller-Hansen, and Jan C. Minx**, “Dataset on the Adoption of Historical Technologies Informs the Scale-up of Emerging Carbon Dioxide Removal Measures | Communications Earth & Environment,” *Communications Earth & Environment*, October 2023, *4* (1), 397.

**Penisa, Xaviery N., Michael T. Castro, Jethro Daniel A. Pascasio, Eugene A. Esparcia, Oliver Schmidt, and Joey D. Ocon**, “Projecting the price of lithium-ion NMC battery packs using a multifactor learning curve model,” *Energies*, 2020, *13* (20).

**Petropoulos, Fotios, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K. Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J. Bessa, Jakub Bijak, John E. Boylan, Jethro Browell, Claudio Carnevale, Jennifer L. Castle, Pasquale Cirillo, Michael P. Clements, Clara Cordeiro, Fernando Luiz Cyrino Oliveira, Shari De Baets, Alexander Dokumentov, Joanne Ellison, Piotr Fiszeder, Philip Hans Franses, David T. Frazier, Michael Gilliland, M. Sinan Gönül, Paul Goodwin, Luigi Grossi, Yael Grushka-Cockayne, Mariangela Guidolin, Massimo Guidolin, Ulrich Gunter, Xiaojia Guo, Renato Guseo, Nigel Harvey, David F. Hendry, Ross Hollyman, Tim Januschowski, Jooyoung Jeon, Victor Richmond R. Jose, Yanfei Kang, Anne B. Koehler, Stephan Kolassa, Nikolaos Kourentzes, Sonia Leva, Feng Li, Konstantia Litsiou, Spyros Makridakis, Gael M. Martin, Andrew B. Martinez, Sheik Meeran, Theodore Modis, Konstantinos Nikolopoulos, Dilek Önköl, Alessia Paccagnini, Anastasios Panagiotelis, Ioannis Panapakidis,**

Jose M. Pavía, Manuela Pedio, Diego J. Pedregal, Pierre Pinson, Patrícia Ramos, David E. Rapach, J. James Reade, Bahman Rostami-Tabar, Michał Rubaszek, Georgios Sermpinis, Han Lin Shang, Evangelos Spiliotis, Aris A. Syntetos, Priyanga Dilini Talagala, Thiyanga S. Talagala, Len Tashman, Dimitrios Thomakos, Thordis Thorarinsdottir, Ezio Todini, Juan Ramón Trapero Arenas, Xiaoqian Wang, Robert L. Winkler, Alisa Yusupova, and Florian Ziel, “Forecasting: Theory and Practice,” *International Journal of Forecasting*, July 2022, *38* (3), 705–871.

**Rypdal, K**, “Empirical growth models for the renewable energy sector,” *Advances in Geosciences*, 2018, *45*, 35–44.

**Smil, Vaclav**, “Examining Energy Transitions: A Dozen Insights Based on Performance,” *Energy Research & Social Science*, December 2016, *22*, 194–197.

UNFCCC, “The Paris Agreement,” 2015.

**Vinichenko, Vadim, Aleh Cherp, and Jessica Jewell**, “Historical Precedents and Feasibility of Rapid Coal and Gas Decline Required for the 1.5°C Target,” *One Earth*, October 2021, *4* (10), 1477–1490.

**Way, Rupert, Matthew C. Ives, Penny Mealy, and J. Doayne Farmer**, “Empirically Grounded Technology Forecasts and the Energy Transition,” *Joule*, September 2022, *6*, 1–26.

## Chapter IV

# Will national renewable costs continue declining?

C. LENNART BAUMGÄRTNER, J. DOYNE FARMER

Since the 1980s, the global average of utility-scale solar photovoltaic (PV) cost has decreased by about two orders of magnitude while the global average onshore wind cost has decreased by about one order of magnitude. But global averages are only part of the story: For both solar and wind, costs vary between countries by about one order of magnitude. We curate a comprehensive database of national costs and build predictive models for future costs by decomposing the levelized cost of electricity (LCOE) into its components and modeling them separately. For solar we find that since 1990 the module cost and the balance of system (BOS) cost have both declined roughly exponentially at rates of 12% per year. In contrast, wind turbine cost has declined roughly exponentially by 4% per year, while BOS cost has not declined at all. This suggests that total global wind cost will approach a floor cost of about 35 USD/MWh, reaching about 43-52 USD/MWh in 2050, whereas global solar cost will continue to decline exponentially, reaching about 3-15 USD/MWh in 2050 and continuing to drop thereafter. For solar, around half of the cross-sectional variations in national LCOE costs is due to variations in capacity factors, though the cost of capital also plays an important role. For wind, the cross-sectional variations are split approximately equally between capacity factor, cost of capital, and investment costs. National BOS costs revert to global costs with a timescale of about 5-7 years, while variations in other factors are more persistent. We develop a simple model for predicting the deviations from global costs and show that it makes reasonable predictions with predictable errors. Finally, we compare our predictions to recent projections of integrated assessment models.

---

**Acknowledgments:** For helpful comments and feedback, we thank François Lafond, Benjamin Wagenvoort, Brendon Tankwa, Cameron Hepburn, and seminar and conference participants at the seventeenth annual meeting of the Integrated Assessment Modelling Consortium (IAMC), 2024.

## IV.1 Introduction

The rapid decline of the costs of utility-scale solar photovoltaic and onshore wind energy over the last sixty years or so has changed views about the feasibility of the energy transformation, injecting a rare note of optimism into the climate debate (Way et al., 2022). In hindsight, this could have been predicted earlier by extrapolating past trends; a few did this, but their voices received limited attention (Farmer and Makhijani, 2010; Farmer and Lafond, 2016). Extrapolation of past trends need to be treated with care (Baumgärtner et al., 2024). After all, trends can change and continuity is not guaranteed. Nevertheless, forecasting future technology costs based on past trends can be useful and has been more accurate than alternatives suggesting future structural breaks (Way et al., 2022). Paying more attention to earlier renewable forecasts could have made it possible to anticipate which technologies were likely to be cheap and plan and improve investment and policy targets (Creutzig et al., 2017; Bolinger et al., 2022b).

At the same time, national costs of renewables vary widely and it is important for national planning and policy to take this into account (Lutz et al., 2015). As we show here, the deviations across countries can be understood by decomposing the costs of renewables into their components. This turns out to also have implications for global costs; as we will see, extrapolating the revealed trends suggests that wind power will drop in cost more slowly and reach a floor, whereas solar will continue to drop exponentially for some time. This suggests that twenty years from now, solar energy will be substantially cheaper than wind energy.

The standard way to compare the costs of capital-intensive renewable technologies with those of more operationally expensive technologies such as fossil fuels is based on the *levelized cost of electricity* (LCOE), which is defined as the average annualized cost of generating one unit of electricity (IRENA, 2024; Bolinger et al., 2022b). The LCOE has substantial limitations, such as excluding grid planning and integration costs (Wang et al., 2022) but remains a convenient way to compare costs. The presented study of LCOE

should therefore be seen as an input to wider energy system models that include omitted system costs and enable forecasts for national energy system and electricity costs.

Here we decompose the formula for LCOE into its different components and develop a forecasting model for each component separately. This allows us to forecast national renewable electricity costs in the absence of future structural breaks and better understand the drivers that determine them.

## IV.2 Decomposition of the LCOE

We consider the LCOE  $\mathcal{L}_j$  of country  $j$ . Over its lifetime an asset produces electricity, which we measure in kWh, and incurs costs measured here in 2020 US dollars. The LCOE translates the initial and ongoing costs for renewable electricity into an effective ongoing cost per year. We follow the standard approach of [Bolinger et al. \(2022b\)](#) and [IRENA \(2024\)](#) and decompose the LCOE of an asset  $i$  in country  $j$  and year  $t$  as

$$\mathcal{L}_{ijt} \equiv \frac{\mathcal{I}_{ijt} \cdot \mathcal{C}_{ijt} + \mathcal{O}_{ijt}}{\mathcal{Y}_{ijt}} \quad (\text{USD} / \text{kWh}). \quad (\text{IV.1})$$

As explained in more detail below, the components of  $\mathcal{L}_{ijt}$  are the total investment cost  $\mathcal{I}_{ijt}$  (USD/kW), the capital recovery factor  $\mathcal{C}_{ijt}$  (%/year), the annual operating and maintenance cost  $\mathcal{O}_{ijt}$  (USD/kW/year), and the capacity factor  $\mathcal{Y}_{ijt}$  (kWh/kW/year), also called annual electricity yield) (Table B.1 in the Appendix summarizes our notation.) While ideally this analysis should be done at the level of individual electricity-generating assets, since data is not available at this level, we are forced to operate at the country level. For convenience, we omit the indices  $ijt$  for the rest of this paper. We use the national average components  $(\mathcal{I}, \mathcal{C}, \mathcal{O}, \mathcal{Y})$  to forecast the national average LCOE  $\mathcal{L}$ . While (IV.1) does not hold exactly for the national averages, we show in Appendix B.3.1 that it is a good approximation.

The *investment cost*  $\mathcal{I}$  includes both equipment costs and soft costs, such as installation and permitting costs, that occur before electricity production begins. It is typically measured in cost (USD) per electric capacity (kW). The investment cost can be further

divided into the equipment cost and the Balance of System cost (BOS). The largest piece of equipment for solar energy is the photovoltaic module, and for wind, it is the turbine; in either case we define the cost as  $\mathcal{M}$ . The Balance of System (BOS) cost  $\mathcal{J}$  aggregates all other equipment and soft costs included in the total investment cost  $I$ , so that we can decompose the investment cost as

$$\mathcal{I} = \mathcal{M} + \mathcal{J} \tag{IV.2}$$

for both solar and wind.

The *capital recovery factor*  $\mathcal{C}$  depends on the assets' lifetime and the weighted average cost of capital (WACC). After an initial investment  $\mathcal{I}$ , the asset operator makes annual payments to its investors (who may or may not be the same entity as the asset operator). As a result, investors recover their initial investment plus a premium.  $\mathcal{I} \cdot \mathcal{C}$  is the annual recovery of the original investment, expressed in USD/kW/year.

The *operating and maintenance cost*  $\mathcal{O}$  is the annual expense of operating the renewable asset, averaged over the asset's lifetime. For renewables, it is typically measured in USD/kW/year (as opposed to USD/kWh for fossil fuels) since there is little to no variable cost that depends on fuel consumption.

The *capacity factor*  $\mathcal{Y}_j$  is the annual electricity generation (kWh) per electric capacity (kW). It accounts for variable environmental characteristics, such as sunshine hours and average wind speed, and also includes asset downtime due to maintenance and other operational constraints.

As evident from Equation (IV.1), the LCOE  $\mathcal{L}$  goes to zero when  $\mathcal{I}$  and  $\mathcal{O}$  go to zero. In that sense, the capital recovery factor  $\mathcal{C}$  and capacity factor  $\mathcal{Y}$  are modifiers to the investment and operating costs. As we will show,  $\mathcal{C}$  and  $\mathcal{Y}$  can nevertheless drive substantial variation in national LCOEs.

The lifetime O&M costs  $T \cdot \mathcal{O}$  (where  $T$  is the renewable asset's lifetime) is 25-75% of the (overnight) total investment cost  $\mathcal{I}$ . When accounting for additional investor premia

for  $\mathcal{I}$ , the investment cost is around one order of magnitude larger than the lifetime O&M cost. Additionally, there is little granular data on O&M costs. As discussed later, we therefore take a simplified approach to forecasting  $\mathcal{O}$ ; more accurate LCOE forecasts warrant more extensive data collection for  $\mathcal{O}$  in the future.

The most important components of the LCOE for wind and solar are shown in Figure IV.1, which makes key properties of the components of the LCOE apparent from visual inspection alone. (See also the summary in Table IV.1).

- In panel 1(a), we see that solar investment costs have dropped by almost three orders of magnitude since 1975, at an average rate of 12% per year. Since 2000, when data on BOS costs first became available, the module cost and BOS cost of solar have been comparable, and have dropped at similar rates. National costs vary substantially, with a standard deviation that is currently about 306 USD/kW, which has also dropped roughly exponentially over time (Elshurafa et al., 2018; Nemet et al., 2020; Nemet, 2006).
- In panel 1(b), wind investment costs have also dropped, though at a much slower average rate of about 3% per year. Together with declines in the cost of capital, this is the main driver of LCOE declines. In 1980 wind was more than an order of magnitude cheaper than solar; it is now, on average, slightly more expensive. The standard deviation of wind investment costs is about 437 USD/kW and has not declined. While turbine costs have declined by an average of 4% per year since 1980, they vary based on origin, with Chinese-built turbines currently about a factor of 3 cheaper than western-built turbines. Most importantly, while BOS costs are highly variable, with a standard deviation of about 400 USD/kW, over a span of almost three decades, *average BOS costs have not declined* (Key et al., 2022; Lacal-Arántegui et al., 2018).
- In panels 1(c) and 1(d), the cross-sectional standard deviation of the capacity factors varies by about 0.5 MWh/kW for wind and solar. The average has remained roughly

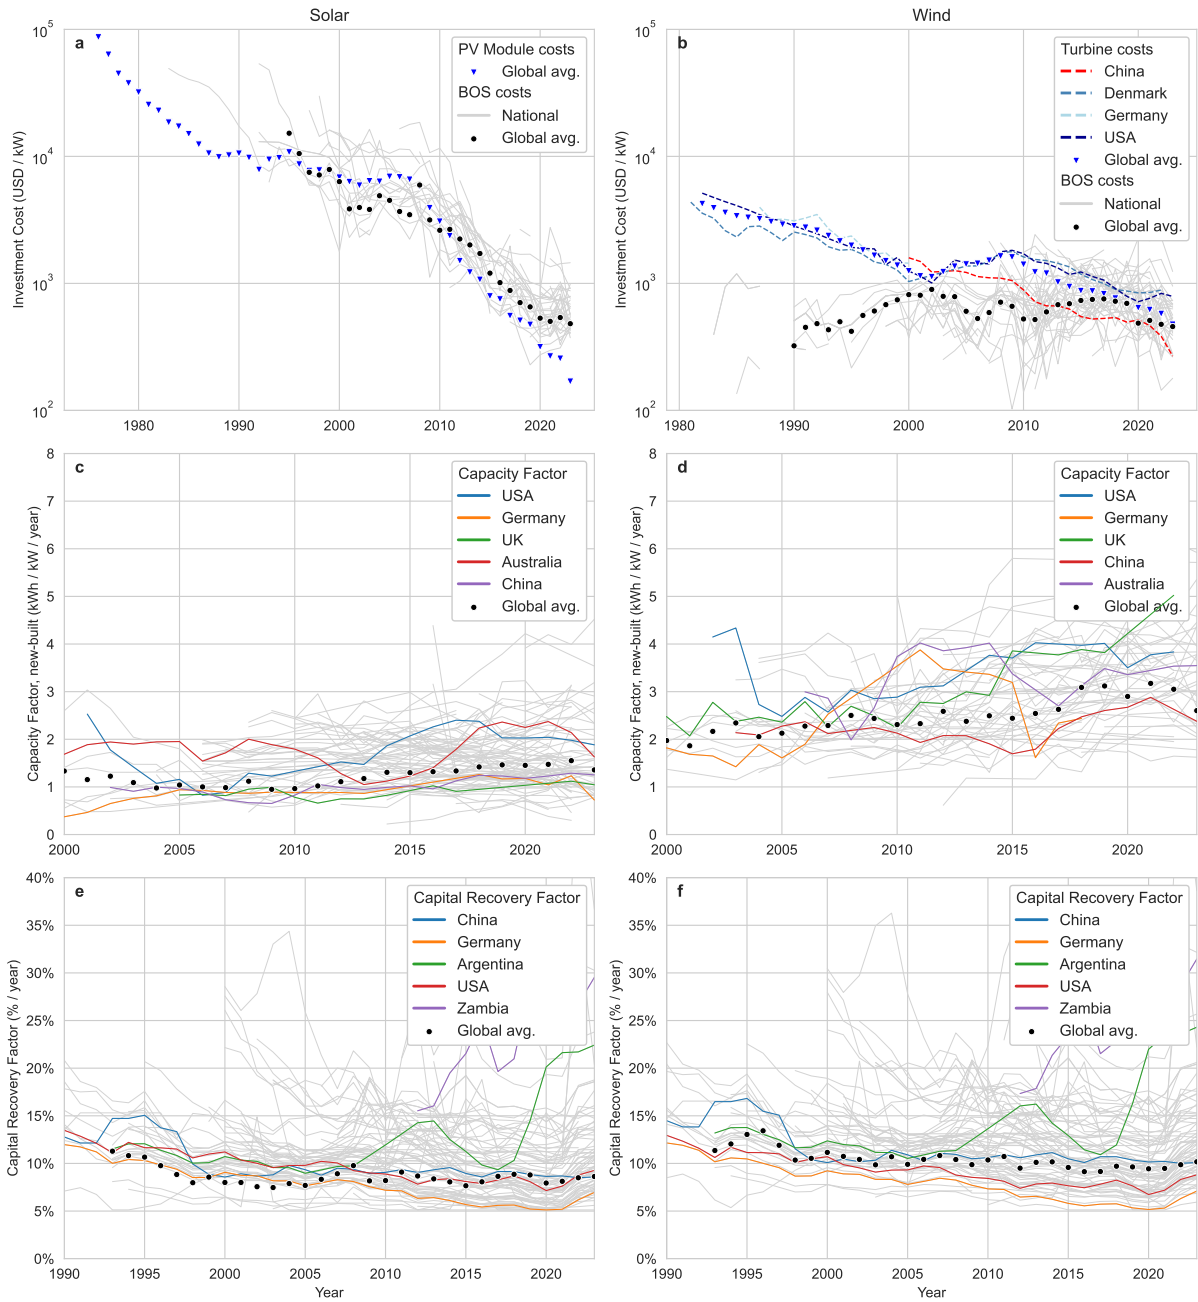


Figure IV.1: **Historical evolution of solar (left) and wind (right) LCOE components.** (a, b) Investment cost components over time. Gray lines show national BOS costs. Black dots indicate the global average BOS cost. Blue triangles show the global average PV module costs for solar and turbine costs for wind. For wind, we see that turbines manufactured in China have been cheaper than western-manufactured alternatives after around 2005. All cost data in this paper is inflation- and currency-adjusted to 2020 USD. (c, d) National capacity factors over time. We observe large variations and temporal fluctuations for national capacity factors driven mostly by external weather factors. Wind has higher capacity factors than solar in most countries. The global average, indicated by black dots, appears to be increasing in both technologies. This is driven primarily by additional countries entering the dataset, not by a long-term increase in national capacity factors. e, f National capital recovery factors (CRFs) over time. CRFs are very similar for solar and wind since they both depend largely on the national financial market. Changing market conditions over time and between countries cause large fluctuations and cross-sectional variation. (We neglect  $\mathcal{O}$  because we were unable to find good data, and in any case it is small)

constant for solar and increased slightly for wind. However, it is important to note that capacity factors are not realized values but rather are estimates. Thus trends may reflect changes in weather, unplanned downtime or forecasting methodology rather than real increases in lifetime capacity factors (Boccard, 2009; Boretti and Castelletto, 2020). (If early projects occupy the most favorable sites, one would expect capacity to decrease, but this is the opposite of what is observed).

- In panels 1(e) and 1(f) we see that capital recovery factors vary substantially, with a cross-sectional standard deviation of about 3.4% for solar and 5.7% for wind. This variation is mostly driven by interest rates, which are generally higher and more volatile in less developed countries (Egli et al., 2018). As a result, capital recovery factors are generally lower for developed countries and higher for less developed countries where investors require higher returns on their investments (Mulugetta et al., 2022). Capital recovery factors have declined continuously by around 1% annually in most countries between the 1990s and 2020 but can be subject to rapid changes due to other macroeconomic dynamics. The high volatility in Argentina’s interest rates, or the sharp increase in interest rates in 2020 following government default, illustrates this.

Table IV.1 presents typical values for each component of LCOE and each technology in 2023. We see that BOS costs are the dominant component of the investment cost for solar energy, as much as 3.5 times module costs, whereas for wind, BOS costs and turbine costs are comparable. BOS costs vary by as much as a factor of 7.5 for wind and a factor of 3 for solar. The capital recovery factor varies by as much as a factor of 4 due to differences in the financial environment. The capacity factor varies by more than a factor of 2 for both solar and wind. For operating and maintenance costs, we cannot make a comparable statement due to the data limitations (Steffen et al., 2020). However, as we will discuss later, the comparably small size of  $\mathcal{O}$  means that we can nevertheless explain almost all of the cross-sectional variance. In summary, all the factors of LCOE except for solar module

Technology	Component	Cost range	Model
Solar	Module cost $\mathcal{M}$	260 USD / kW	Global forecast for costs decline based on Wright's law, conditional on global Solar deployment
Solar	BOS cost $\mathcal{J}$	300 - 900 USD / kW	Global average BOS cost declines are forecast based on (global) Wright's law and national deviations are forecast using a mean-reverting process
Wind	Turbine cost $\mathcal{M}$	400 - 900 USD / kW	Regional forecast for cost decline in either Western or Chinese turbines, based on Wright's law, conditional on manufacturing location and global wind deployment
Wind	BOS cost $\mathcal{J}$	200 - 1500 USD / kW	No long-term trend. Both global average BOS costs and national deviations are based on mean-reverting processes
Solar Wind	Capital recovery factor $\mathcal{C}$	5 - 20 %/year	We assume the asset lifetime and the weighted average cost of capital (WACC), which depends on future risk-free rates and technology-specific markups, are given.
Solar Wind	Operating & Maintenance $\mathcal{O}$	$\mathcal{I} \cdot 1\%$ /year (Solar), $\mathcal{I} \cdot 3\%$ /year (Wind)	Constant forecast relative to the total investment cost $\mathcal{I}$ , no long-term trend relative to $\mathcal{I}_j$
Solar Wind	Capacity Factor $\mathcal{Y}$	10 - 25 % (Solar), 25 - 45 % (Wind) of 8240 hours / year	No long-term trend. We use an average for each country and technology, which depends on geography and the local electricity system.

Table IV.1: **Overview of Solar and Wind component models.** Typical component values are given for the year 2023. The component models for solar and wind differ based on the available data and empirically tested differences. For example, future solar BOS costs are forecast using Wright's law and a model for the national deviations, while future wind BOS costs are forecast using a mean-reverting process.

cost vary widely. BOS costs, capital recovery factors and capacity factors are all important determinants of the variability of national costs.

### IV.3 Forecasting the components of LCOE

To forecast the components of renewable technology costs, we follow a statistical approach and describe the quantitative properties of our component time-series. These statistical properties can then be used to predict future component values if historical trends continue. As in many Machine Learning examples, this approach omits the crucial topic of causal relationships and future structural breaks. Nevertheless, statistical forecasts can be useful in practical applications (see [Lafond \(2025\)](#) for a detailed discussion of this topic).

Following this statistical approach, [Nagy et al. \(2013\)](#), [Farmer and Lafond \(2016\)](#), and [Lafond et al. \(2018\)](#) show that technology costs can be forecasted independently. As discussed in [Lafond \(2025\)](#), considering technology dependence is an important avenue for future research. Technologies can depend on each other (e.g., supply chains or funding), so considering correlations may significantly improve forecasts in the future. The fact that the components of renewable technology costs have followed different trends suggests that it is also useful to forecast each of them separately and then combine the forecasts. This is helpful for both national and global forecasting and helps to understand why national costs vary so much. As described in [Appendix B.3.6](#), we have found no strong evidence of correlation between LCOE components outside the joint trend of solar modules and BOS but this may change if more data is collected.

For this purpose, it is convenient to work with logarithms. Since operating and maintenance costs for renewables are relatively small and data is largely unavailable, we assume that they are a constant proportion of total investment costs,  $\mathcal{O} = \gamma \cdot \mathcal{I}$ , where  $\gamma \approx 1\%$  for solar and  $3\%$  for wind ([Riva et al., 2018](#); [Bolinger et al., 2022a](#); [Wiser et al., 2019](#)). Substituting and taking the log of Equation (IV.1) yields the additive expression

$$\log \mathcal{L} \approx \log \mathcal{I} + \log(\mathcal{C} + \gamma) - \log \mathcal{Y}. \quad (\text{IV.3})$$

[Figure IV.1](#) suggests that these terms behave differently through time and that the behavior of wind differs from that of solar, so that there is value in considering them separately, and in decomposing the investment cost  $\mathcal{I}$  into its components  $\mathcal{M}$  and  $\mathcal{J}$ .

Moore’s and Wright’s laws provide the two best methods discovered so far to forecast average global costs based on historical data ([Nagy et al., 2013](#); [Farmer and Lafond, 2016](#); [Lafond et al., 2018](#); [Meng et al., 2021](#)). Moore’s Law predicts that costs will decline roughly exponentially over time, while Wright’s Law predicts a rough power law relationship between costs and cumulative deployment of a technology, where the word “rough” is used to indicate that these relationships are noisy. Moore’s and Wright’s laws have been tested

on about 60 technologies and shown to provide better forecasts than several alternatives (Nagy et al., 2013; Meng et al., 2021; Farmer and Lafond, 2016; Lafond et al., 2018). Nagy et al. (2013) showed that Moore’s law and Wright’s law forecasts are similar because in all the data sets studied so far, production increases roughly exponentially while costs decline roughly exponentially, which makes Wright’s and Moore’s laws equivalent. This will likely change once wind and solar become mature technologies and the rapid exponential growth of their deployment slows down. In this case, their forecasts will differ, but with our current state of knowledge, the best method of forecasting remains an open question (Lafond et al., 2020).

If we assume that the trends observed in Figure IV.1 persist, then in the future both module costs and BOS costs for solar will drop roughly exponentially according to Moore’s law, as they have in the past. In contrast, wind turbine costs will drop roughly exponentially, but its BOS costs will remain roughly constant. This implies that wind will reach a floor cost of 35 USD/MWh. The resulting forecasts for solar and wind are compared in Figure IV.2. We forecast that solar costs will continue to drop roughly exponentially, reaching a cost of about 3 USD/MWh in 2050, whereas the rate of decline in overall wind costs slows down, reaching a cost of 43 USD/MWh in 2050. Thus, in 2050, solar energy will become around an order of magnitude cheaper than wind energy. We could also have used Wright’s Law, but then we would have needed to postulate a scenario for future production. If we had assumed that deployment continues at current roughly exponential rates, then the forecasts would have been the same as they are under Moore’s law. This is unlikely because under current rates of exponential deployment solar and wind would become the sole providers of useful energy around 2036 due to the large efficient gains in a fully electrified energy system (Baumgärtner et al., 2024; Way et al., 2022). Thus, before this point is reached, deployment must slow down. In this case the forecast of the two methods differ, and under Wright’s law we expect their costs in 2050 to be somewhat higher than predicted above. As shown in the Appendix, under plausible decarbonization scenarios from Way et al. (2022), the difference has little effect on wind,




Figure IV.2: **Historical evolution and forecast of global average solar (red) and wind (blue) LCOE.** Dots show the historical evolution of global average solar and wind LCOE. The solid line indicates the historical long-term trend based on decreasing PV modules, wind turbines, and solar BOS costs. The dotted line indicates our long-term forecast for each technology, with shaded areas indicating the 50% and 90% confidence intervals. For wind, we forecast a floor cost of around 35 USD/MWh due to the constant long-term trend in BOS costs.

and increases median solar costs to 12-15 USD/MWh in 2050.

We want to stress that, although it has become standard in the literature to use Wright’s law to forecast technological change, there are strong dissenters in favor of Moore’s law (Funk and Magee, 2015); a study of technological change during World War II indicates that stimulating demand causes prices to drop under Wright’s law, but it also shows that Wright’s law and an overall exponential trend have roughly the same predictive power (Lafond et al., 2020). Hence here for simplicity we use Moore’s law, and present more detailed evidence using Wright’s law in the appendix.

The trends we observe here are plausible for more fundamental reasons. The use of

larger turbines and higher hub heights has been a major factor in making wind energy cheaper, next to reduced material need, manufacturing scale, and labor productivity (Elia et al., 2020). The lack of decline in BOS costs for wind observed here suggests that as turbines become larger, they also become more difficult to install, offsetting cost reductions from learning (Key et al., 2022; Lacal-Aránategui et al., 2018). Solar photovoltaic panels, in contrast, have maintained a similar size and there are no such confounding factors, so as solar farms become larger, a combination of economies of scale and learning make the installation costs and possibly the other components of BOS costs drop. It is nonetheless remarkable that module costs and BOS costs for solar have decreased at such similar rates.

To assign errors to these forecasts, we use the method developed by Farmer and Lafond (2016) which was originally applied to module costs alone but has also been tested in forecasting the costs for many technologies. This method takes into account errors caused by parameter estimation and the noisy fluctuations causing deviations from pure exponential behavior. It has been tested on more than 50 time series and has been shown to do a good job of predicting forecasting errors. It has so far produced good predictions of future renewable costs, with forecasts typically within 95% confidence intervals (Way et al., 2022). Here, we assume that the capacity factor and cost of capital remain constant on average, which is plausible since capacity factor is largely determined by geography and the cost of capital is largely determined by the financial climate in each country. Our forecasts can be used to determine the impact of different exogenous financial conditions on local renewable costs. This is particularly important for developing countries, where high financing costs can either hinder renewable adoption or foreign financing mechanisms may be used to accelerate a clean electricity grid (Lonergan et al., 2023; Polzin et al., 2021; Mulugetta et al., 2022; Komendantova et al., 2019).

As a reality check, we regress the components of the logarithm of the LCOE on the right side of Equation (IV.3) against the total LCOE (which is obtained from other sources but subject to data cleaning). The regression explains about 99% of the variance in the national LCOE data, and the regression coefficients are all within a 95% confidence interval

of one, indicating that the data are fairly consistent (see Section B.3.1 in the Appendix). The fact that we do not explain 100% of the variance is because there are some non-zero residuals, which have a heavy-tailed distribution, indicating data errors. Thus, while the data are not perfectly consistent, the errors caused by the inconsistencies are relatively small in relation to the expected forecasting errors.

To provide some insights into the variation in national prices at a given time we study the cross-sectional deviations of the national costs from the global average. It is more convenient to work in terms of logarithms, so we define the variables  $\lambda_{jt} = \log \mathcal{L}_{jt}$ ,  $\iota_{jt} = \log \mathcal{I}_{jt}$ ,  $\kappa_{jt} = \log(\mathcal{C}_{jt} + \gamma)$  and  $v_{jt} = \log \mathcal{Y}_{jt}$ , which implies that  $\lambda_{jt} = \iota_{jt} + \kappa_{jt} + v_{jt}$ . Taking the variance of both sides gives

$$\text{Var}(\lambda_{jt}) = \text{Var}(\iota_{jt}) + \text{Var}(\kappa_{jt}) + \text{Var}(v_{jt}) + \text{cross covariance terms}, \quad (\text{IV.4})$$

or in words, the variance of the logarithm of the LCOE is the sum of the variance of investment costs plus the variance of the capacity factor plus the variance of the capital costs, plus cross-covariance terms. Because the individual terms of the LCOE decomposition depend on very different factors, we expect the cross-covariance terms to be close to zero. The cost of capital is determined by macroeconomics and politics, the capacity factor is largely determined by climate and geography, and the investment cost is determined by technology and installation costs. This is confirmed: When we estimate the cross-covariance terms we find that they are all within two standard deviations of zero. We can compute the fractional contributions of each term by dividing each term on the right by  $\text{Var}(\lambda_{jt})$ , see Table IV.2. This implies that for solar about 50% of the variance in cost is due to the capacity factor, 25% due to investment cost and 25% due to the cost of capital. In contrast, for wind only 30% is due to the capacity factor, 30% is due to investment cost and about 40% is due to the cost of capital.

*Solar (2022)*

Component	Cross-sectional mean	log $\mathcal{L}$ share	Variance share
$\mathcal{L}$	97 USD/MWh	(100%)	(100%)
$\mathcal{I}$	985 USD/kW	42%	25%
$\mathcal{Y}$	1.2 MWh/kW	43%	50%
$\mathcal{C} + \gamma$	8% + 1%	15%	25%

*Wind (2022)*

Component	Cross-sectional mean	log $\mathcal{L}$ share	Variance share
$\mathcal{L}$	112 USD/MWh	(100%)	(100%)
$\mathcal{I}$	1440 USD/kW	43%	32%
$\mathcal{Y}$	2.1 MWh/kW	45%	27%
$\mathcal{C} + \gamma$	9% + 3%	12%	41%

Table IV.2: **Cross-sectional standard deviation and variance of LCOE and its components.** For the year 2022, we compute the cross-sectional mean and standard deviation of the LCOE and its components (unweighted). We can further use Equation (IV.4) to estimate the share of the variance that each component contributes to the variance of the logarithm of the LCOE. We leave out the cross-variance terms, contributing less than 10% for solar and less than 5% for wind.

#### IV.4 Forecasting national costs

To construct a model for predicting national costs based on historical data, we build on previous work in forecasting global costs (Farmer and Lafond, 2016; Lafond et al., 2018). We use already well-tested methods to make global forecasts and construct new models for forecasting national deviations from global forecasts. For each component of the LCOE, we pool the data for each country and build a single model for all countries. Our approach is in contrast to previous studies that built separate models for each country (Elshurafa et al., 2018; Helveston et al., 2022; Huenteler et al., 2016; Kothari et al., 2023). Building separate models for each country means that ratio of parameters to data is high, which causes overfitting. To make optimal forecasts one needs to strike the right balance in the bias-variance tradeoff. Bias refers to whether or not a model is capable of fitting the data, and variance refers to errors in parameter estimation, which result in overfitting. In situations like this where data is limited, variance is usually the dominant source of error,

and it is essential to keep models simple (Baumgärtner et al., 2024).

We forecast each component of the LCOE as follows (see also Table IV.1).

*Module costs.* We take advantage of the fact that module costs are essentially global (Masson and Kaizuka, 2022) and make use of methods developed in previous studies (Farmer and Lafond, 2016; Lafond et al., 2018), as discussed above. Since these have been validated by making out-of-sample forecasts for many different technologies, we do not need to test them here. There are a few exceptions, such as Japan, which produces and sells a majority of its solar panels domestically (Raupach-Sumiya et al., 2015). We capture the associated costs in the BOS costs through equation (IV.2) but cannot differentiate the effects of domestic panel production from other effects increasing national BOS costs.

*Turbine costs.* Wind turbines are manufactured primarily in the West (US & Europe) and China, with Chinese turbines now costing around one-third of Western turbines (IRENA, 2024). We can use either the previous global forecasts (Farmer and Lafond, 2016; Lafond et al., 2018) for the average cost of Western vs Chinese turbines or differentiate between the two locations and forecast them separately. However, we cannot identify the preferred method due to low statistical significance of any comparison. For simplicity, we apply the global forecasts here.

*Balance of system costs.* Forecasting BOS costs is complicated by the fact that BOS costs behave very differently for solar and wind. As noted, solar BOS costs closely track module costs, whereas wind BOS costs fluctuate, but the global average is roughly flat. This means that we need to take somewhat different approaches for solar and wind.

We begin with solar. As seen in Figure IV.1, solar energy national BOS costs closely track global BOS costs. While the cross-sectional variation is substantial, it is much smaller than the variation in time, which spans almost two orders of magnitude. This makes the national BOS costs highly correlated, as established in Appendix B.3.4. To cope with this we first take logarithms and then decompose the national costs into the global average and the national deviations from the global average. Since the global average behaves very similarly to the module costs, we use the same methods to forecast it as

noted above.

For wind, national costs are similarly correlated with each other, as established in Appendix B.3.5. Cross-sectional variation is substantial, but individual countries' deviation from the global average are mostly temporary. We can, therefore, take a similar approach to solar and decompose global and national logarithms of the BOS costs. However, unlike solar, the global average does not show a significant long-term trend. We therefore use a mean-reverting process to forecast the average, similar to McNerney et al. (2011) that forecast fossil-fuel based technology costs.

Because they are essentially stationary, to forecast the national deviations of the BOS costs for solar or wind, or the global BOS costs for wind, we assume an AR(1) model of the form

$$x_{j,t+1} = \rho x_{jt} + \sigma n_{jt}, \quad (\text{IV.5})$$

where  $\rho$  is a reversion parameter, with  $\rho \approx 0.85$  for solar deviations,  $\rho \approx 0.7$  for wind deviations and  $\rho \approx 0.5$  for global wind costs.  $n_{jt} \sim N(0, 1)$  is Gaussian noise with mean zero and standard deviation one and  $\sigma$  controls the noise level. For both solar and wind  $x_t$  is the logarithmic deviation of the BOS cost at time  $t$ , i.e.  $x_{jt} = \log \mathcal{J}_{jt} - \frac{1}{N_t} \sum_{j=1}^{N_t} \log \mathcal{J}_{jt}$ . Note that there are no country-specific parameters. This means that all countries revert to the same global mean, which is roughly what we observe. This assumptions reflects the correlation findings described above, although it is not a unique choice. Therefore, we validate the model in Appendix B.3.4 and B.3.5 using out-of-sample testing.

*Capacity factors* are given mostly by geography, so they can be taken as country-specific constants, corresponding to their historical average for each country. We can either estimate the parameter based on historical electricity generation (subject to noise from annual weather patterns) or ex-ante estimates by project developers (subject to noise and biases from different estimation methods). For wind, we estimate the capacity factors based on historical electricity generation since ex-ante estimation biases can be large. For solar, there is little difference between both methods. Here, we use ex-ante estimates by project developers for slightly higher out-of-sample accuracy. We ignore the rising trend

of capacity factors as we believe it is spurious.

*Capital recovery costs* are determined by the national government bond yields. The best way to treat this depends on the purpose of the forecast. For someone who wants to know hypothetically how costs might change if bond yields change, it is appropriate to make conditional forecasts treating this as an exogenous variable. However, for our purposes here we want to make unconditional forecasts, so we assume bond yields to remain constant at their current national levels.

To test our models, we divide the data into training sets (also called the in-sample data) and test sets (also called the out-of-sample data). We fit the parameters of our model on the training set and validate the model on the test set. We take 5-year moving windows as training sets (e.g., 2000-2004) and the remaining data in the future relative to this point as the test set (e.g., 2005-2024). We do not attempt to validate the global components of the model, as we have only two very short time series, and any tests would be meaningless, and the global models have already been tested on many technologies in previous work. In Figure IV.3 we compare predicted vs. actual values for the deviations of the BOS costs from the global average. As expected given our use of an AR model, the correlations decay exponentially with a time constant of  $1/(1 - \rho) = 6.6$  years for solar and 5 years for wind. This is because national deviations are not persistent – they revert to the global mean. A plausible interpretation is that this is due to learning effects across countries.

To quantify this, we compute the number of degrees of freedom, which is a measure of the effective number of independent variables in the time series (See definition in the SI). If we use the raw BOS costs for solar we find only 9 degrees of freedom. In contrast, the logarithmic deviations from the national average have about ten degrees of freedom, meaning that tests of forecast performance in the validation set are much more statistically significant. The situation for wind is simpler, due to the fact that for wind the national BOS cost time series have about 17 degrees of freedom, whether or not we subtract the global average.

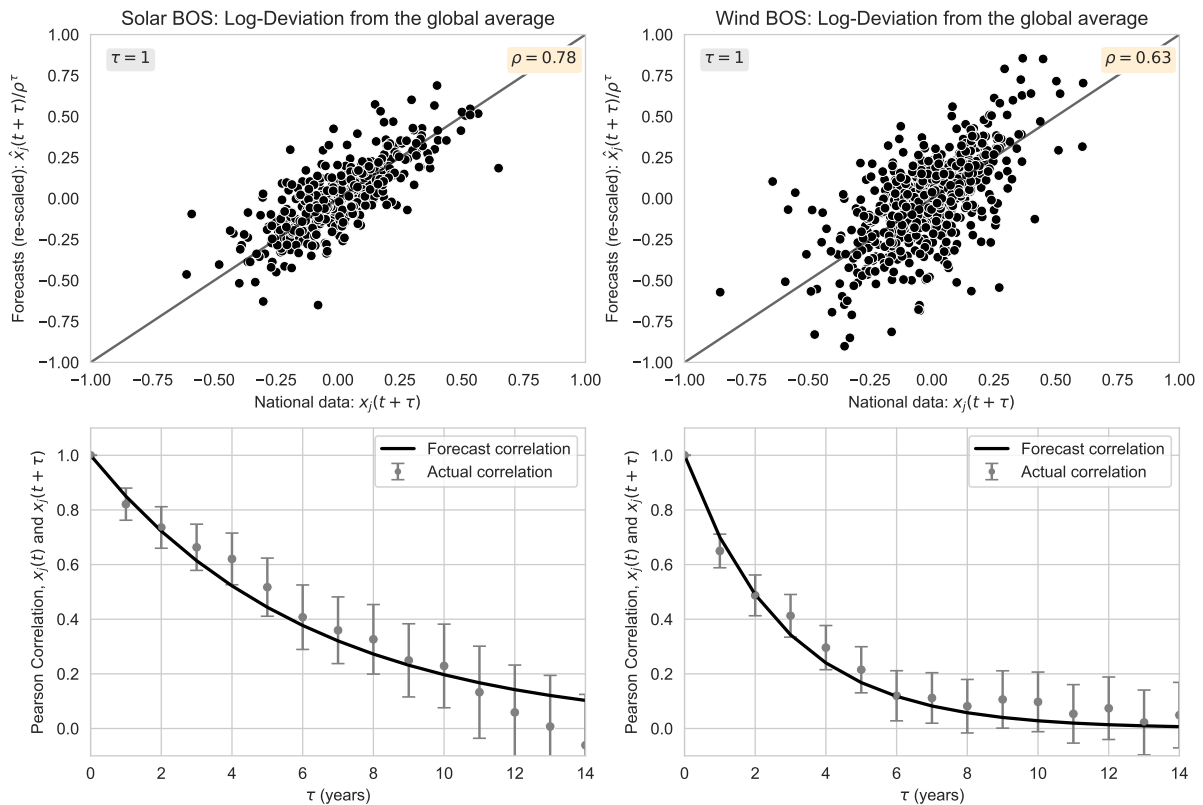


Figure IV.3: **Forecast and actual deviation of national BOS costs from the global average for solar (left) and wind (right).** **Top panels** X-axis shows the national (log-10) deviations from the global average  $x_j$  for times  $t + \tau$ , with  $\tau = 2$ . Y-axis shows the forecast for  $x_j(t + \tau)$  made at times  $t$ ,  $\hat{x}_j(t + \tau)$ , re-scaled to match the cross-sectional variance of the data. Data and forecast are highly correlated with Pearson correlation coefficients of  $\sim 0.78$  for solar and  $\sim 0.63$  for wind. **Bottom panels** Correlation between  $x_j(t)$  and  $x_j(t + \tau)$  with 95% confidence intervals (dots) and  $\hat{x}_j(t + \tau)$  (black line) for different values of  $\tau$ .

## IV.5 Discussion

Our work here shows the value of decomposing costs into their components. Our most important result is that the BOS costs for solar have dropped in tandem with module costs, while the BOS costs of wind have remained roughly constant. This indicates that if statistical trends continue, wind costs will reach a plateau while solar costs will continue to drop exponentially. A plausible explanation for this is that the physical nature of solar panels has hardly changed, and solar farms are easy to build, while the cost declines in wind turbines are substantially due to using bigger turbines, which makes them harder to install.

Our second result is that the difference between national BOS costs tend to revert

to the global mean. This suggests that there is substantial learning transmitted across countries that occurs on a timescale of a few years.

Unlike investment costs or capacity factors, the cost of capital depends on economic factors that are potentially under the control of policymakers in governments, central banks and other financial institutions, or international organizations like the IMF. The cost of capital accounts for about 25% of the variance for solar and 40% for wind, so the differences are substantial. This also means that global changes in interest rates, such as the higher interest rates in recent years, can have a substantial effect on cost competitiveness.

In the past the projections of integrated assessment models for renewable costs were very poor, substantially overestimating future renewable costs. This is in contrast to time series-based forecasts, which have done quite well ([Way et al., 2022](#)). The good news is that, as shown in the Supplementary Material, the projections of integrated assessment models are now much more closely aligned with ours. Nevertheless, some challenges remain, such as keeping scenarios up to date and ensuring that the ensemble of scenarios reflects the uncertainty around future technology costs ([Guivarch et al., 2022](#); [Peters et al., 2023](#)). We hope that the improved forecasts we give here will be useful in honing IAM forecasts further.

These findings come with a number of caveats and limitations. Firstly, as discussed in the introduction, we focus on onshore wind and utility-scale solar photovoltaics, leaving many questions around adjacent technologies (e.g., offshore wind and rooftop solar) unanswered, despite their pivotal role in the energy transition. Some of our findings may apply to them as well but require further data collection and study to confirm. Secondly, our focus on levelized costs of electricity limits the application of our findings. While we hope that our forecasts for LCOE and its components are useful to energy system studies, further analysis of grid integration is required to make direct implications to the cost of the energy transition in different countries. Thirdly, our focus on the statistical description of renewable costs implies that our forecasts will only hold if historical trends

continue into the future. This assumption has held in the past but is not guaranteed to hold indefinitely, particularly in the presence of structural breaks. Furthermore, the statistical nature of this work omits important causal relationships. Further study into the causal factors underlying the statistical relationships is important for investment and policy decisions (Lafond, 2025).

## IV.6 Methods

### IV.6.1 Data collection and cleaning

Validating our forecasts is important to improve the reliability of our forecasts (Baumgärtner et al., 2024). Statistical validation, however, has substantial data requirements. To maximize the available data and develop a widely applicable model, we apply and validate the same methods across a novel sample of 102 countries for solar and 111 countries for onshore wind. Here, we provide a brief overview of the dataset. Additional details are given in Appendix B.2.

We collected and extensively cross-validated national average data for LCOE  $\mathcal{L}_j$ , total investment cost  $\mathcal{I}_j$ , cost of capital, electric capacity, and electricity generation by manually scraping the extended scientific and gray literature. Next to the investment cost shown in Figure IV.1 and LCOE shown in Figure IV.2, Figure IV.4 shows the capacity of installed solar and wind projects for different countries. We see that national capacities have grown by up to 3 orders of magnitude for wind (e.g., South Africa) and 4 orders of magnitude for solar (e.g., Saudi Arabia) in the last two decades alone. So far, there is no statistically significant departure from the exponential trend (Baumgärtner et al., 2024).

Due to the diversity of sources, the data is naturally noisy and subject to inconsistent data-cleaning and aggregation procedures. This particularly applies to the LCOE  $\mathcal{L}_j$ . For example, data providers sometimes fill data gaps are sometimes filled using ad-hoc assumptions. To reduce inconsistencies we perform several corrections on the LCOE data. These corrections are largely modelled on IRENA’s data (IRENA, 2024) since it is the

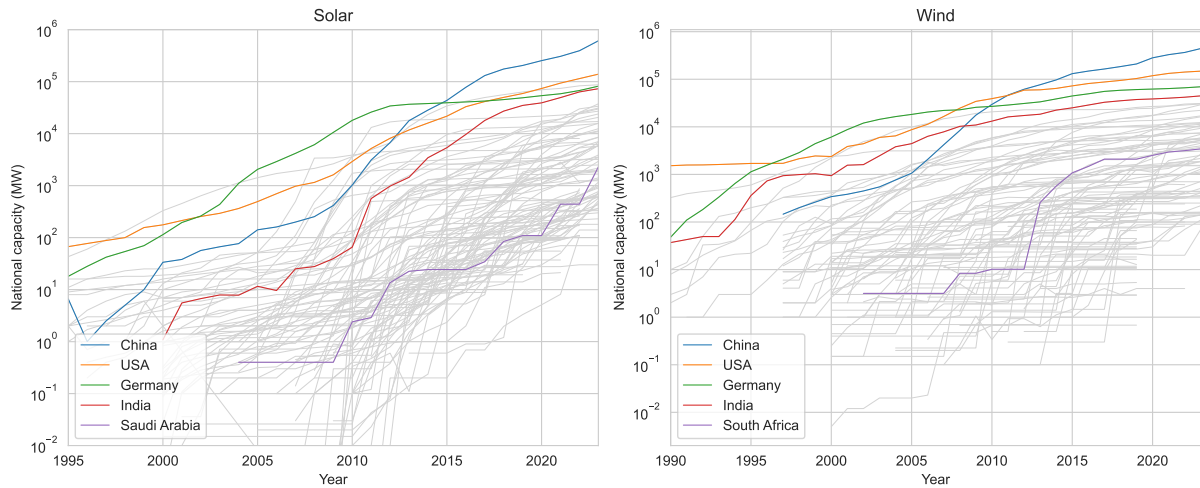


Figure IV.4: **Solar (left) and onshore wind (right) capacities for different countries and years** Over the last 40 years, solar and wind capacities have grown approximately exponentially across the world. While there is significant fluctuation between countries and years, there has been not significant departure from exponential growth so far (Baumgärtner et al., 2024).

largest single source of LCOE data in our dataset and arguably the most prominent source of global and country-level data.

*Cost of capital.* We harmonize the LCOE data by using a consistent cost of capital across time and countries. For data prior to 2021, IRENA applies simplified cost of capital assumptions. After 2021, they use more granular data (IRENA, 2023). We therefore correct the data to obtain a consistently granular cost of capital throughout time and across countries. The impact of this correction on our model is minimal since we take the cost of capital as given.

*Capacity factor.* For solar, IRENA’s capacity factor closely aligns with other sources. For wind, we find significant discrepancies due to different estimation methods. IRENA largely relies on ex-ante estimates by project developers. These can be materially different from ex-post estimates based on the actual electricity generated (see also Appendix B.3.7). To harmonize our data, we use the capacity factor based on electricity generation for wind. This correction can be a limitation to our LCOE forecasts, which is why we validate the capacity factor forecasts independently.

*Operating and maintenance cost.* While IRENA does not provide detailed information on their operating and maintenance cost prior to 2010, we identify some inconsistencies

across the dataset. We harmonize these inconsistencies throughout the data by using our assumptions for the operating and maintenance costs that are based on an independent dataset not used in any other LCOE forecasts. This limits data correction impacts on the LCOE forecasts but does not avoid them completely. However, since the operating and maintenance costs constitute but a small part of the LCOE, the resulting noise is small compared to the variance in other LCOE components.

## IV.6.2 LCOE decomposition

Decomposing the LCOE  $\mathcal{L}_{jt}$  into its components is a non-trivial task due to the nonlinearity of equation (IV.1). As a consequence, Equation (IV.3) does not hold precisely but requires an additional error term,

$$\log \mathcal{L}_{jt} = \log \mathcal{I}_{jt} + \log(\mathcal{C}_{jt} + \gamma) - \log \mathcal{Y}_{jt} + \log \mathcal{E}_{jt}, \quad (\text{IV.6})$$

where  $\log \mathcal{E}_{jt}$  is an additional random variable characterized by an unknown distribution.

We can use our dataset on the LCOE and its components to determine the distribution of  $\log \mathcal{E}_{jt}$ . We find that  $\log \mathcal{E}_{jt}$  does not exhibit significant autocorrelation. Furthermore, the sample mean and sample variance of  $\log \mathcal{E}_{jt}$  do not significantly change over time and across countries. Although some countries and years display notably larger values of  $\log \mathcal{E}_{jt}$ , their difference is not statistically significant due to the small sample size. Consequently we model  $\log \mathcal{E}_{jt}$  as an independent and identically distributed (i.i.d.) random variable, independent of indices  $j$  and  $t$ .

However, we also find that  $\log \mathcal{E}_{jt}$  can take some very large empirical values, indicating that it follows a fat-tailed distribution, such as a Cauchy distribution. This observation prompts further investigation into whether these extreme values result from the nonlinearities inherent in the LCOE calculation or from potential data inconsistencies.

Unfortunately, asset-level data required to study this effect is scarce. A notable exception to this is the dataset provided by [Barbose et al. \(2024\)](#), which consolidates the

investment costs for over 2.4 million solar photovoltaic assets in the United States between 2000 and 2023. Their public dataset covers around 5% of total U.S. solar capacity in the US, although coverage varies substantially by year. Using this data, we see that single assets can cause large and frequent fluctuations in the differences between the weighted average of  $\mathcal{I}_{ijt}$  and the weighted average of  $\log \mathcal{I}_{ijt}$  of up to an order of magnitude. These variations directly translates into large fluctuations in  $\log \mathcal{E}_{jt}$  between consecutive years. For instance, in 2016, the dataset includes one asset with a capacity of 424 MW and exceptionally low installation costs. The asset size is an order of magnitude larger than all other assets in 2015 and 2016 and five orders of magnitude lower in costs than the annual average. Consequently,  $\log \mathcal{E}_{jt}$  fluctuates by an order of magnitude between 2015 and 2016. If we exclude these outliers, extreme values in  $\log \mathcal{E}_{jt}$  vanish. This implies that large values of  $\log \mathcal{E}_{jt}$  are in large parts caused by data issues and we should not use a fat-tailed distribution to model it.

Instead, we discard extreme values of  $\log \mathcal{E}_{jt}$  and adopt a normal distribution to forecast future national LCOE values.

### IV.6.3 Solar module and wind turbine component model

For both solar modules and wind turbines, we employ the experience curve model developed by [Lafond et al. \(2018\)](#) to forecast global average costs, conditioned on the global cumulative deployment of each respective technology.

In the case of solar modules, this approach is particularly straightforward, as the majority of modules are manufactured in China and traded internationally.

For wind turbines, distinctions exist between Western-manufactured and Chinese-manufactured turbines, impacting investment costs and the LCOE. Nonetheless, due to data limitations, we cannot determine with statistical significance whether a locally-specific experience curve model would yield improved forecasts compared to the global alternative. Therefore, we adopt the simpler global model.

We do not independently validate these models, as [Lafond et al. \(2018\)](#) have already

extensively validated the approach across a dataset encompassing over 50 technologies.

#### IV.6.4 Balance of system component model

Through regression analysis, we confirm that solar Balance of System (BOS) costs exhibit a significant experience rate, similar in magnitude to that observed for solar modules. In contrast, we find no significant cost reduction or experience exponent for wind BOS components over the past four decades.

We next examine the cross-correlation structure present within our dataset. Both solar and wind national time series demonstrate significant correlations, with solar showing notably strong correlations, resulting in only nine effective degrees of freedom.

Additionally, we analyze the autocorrelation structure in our dataset, revealing significant autocorrelation at both global average and national scales. This autocorrelation indicates that an autocorrelated forecasting model is suitable for predicting future national deviations from global averages.

For solar, we further compare various forecasting models aimed at predicting future national costs. However, due to substantial cross-sectional correlations, our out-of-sample tests do not yield statistically significant differences between models. Nonetheless, drawing upon theoretical insights from the time-series forecasting literature, we identify a global average model with autocorrelated national deviations as the most appropriate method for projecting future solar BOS costs.

Finally, for both solar and wind technologies, we perform an out-of-sample validation of our stochastic forecasts. To achieve statistical significance, we specifically assess national deviations from global averages and compare forecast percentiles against a theoretical prior derived from surrogate data.

#### IV.6.5 Capacity factor component model

There are two methods to estimate the capacity factor  $\mathcal{Y}_{jt}$ . Using our dataset, there are two options to estimate the trend-stationary capacity factor. 1) We can estimate the

ex-post capacity factor from historical electric capacity and generation. 2) We can use our data for  $\mathcal{L}_{jt}$ ,  $\mathcal{I}_{jt}$ ,  $\mathcal{C}_{jt}$  and  $\gamma$  in Equation (IV.6) to recover a priori capacity factors used in LCOE estimates from the literature. The LCOE-based method yields an ex-ante estimate provided by data suppliers, whereas the electricity-based approach provides an ex-post estimate derived from recent historical generation, serving as a reference for anticipating ex-ante capacity factors.

For solar, the differences between the two estimation methods are minor. We apply the ex-ante capacity factor derived from the LCOE-based approach due to its marginally superior out-of-sample predictive performance.

For wind, the two methods differ substantially. The LCOE-based approach indicates a significant upward trend in capacity factors, whereas the electricity-based method reveals no such trend—a discrepancy previously documented in the literature (Boretti and Castelletto, 2020; Bolinger and Seel, 2018; Bocard, 2009). Given our objective of accurately predicting actual electricity generation relative to installed capacity, we adopt the electricity-based method, treating future capacity factors as constants specific to each country.

#### IV.6.6 LCOE forecasts

Utilizing the decomposition of the LCOE into its components and applying the forecast models developed for each component, we generate forecasts for national-level LCOE values. Additionally, we forecast the global average LCOE, as illustrated in Figure IV.2, based on projections of global average components.

For solar modules and wind turbines, we directly employ global forecasts derived from experience curves. The future costs are contingent upon projected global cumulative deployment levels for each technology. In the forecasts presented in Figure IV.2, we assume an annual deployment growth rate of 30% for solar and 20% for wind. Alternatively, scenario-based deployment projections, such as those outlined by Way et al. (2022), can be applied.

Similarly, solar BOS costs are forecast using the experience curve model. For wind BOS costs, we employ an AR(1) model to project global average costs, resulting in a stochastic floor-cost estimate in the investment cost  $\mathcal{I}_j$  of roughly 450 USD/MW. This corresponds to a floor cost in LCOE  $\mathcal{L}_j$  of roughly 35 USD/MWh.

Regarding the capacity factor and capital recovery factor, we estimate current global averages and assume these values remain constant over time. This assumption is appropriate under the condition that the distribution of future renewable energy deployment mirrors historical patterns. However, the increasing affordability of solar energy may facilitate greater adoption in lower-income regions, potentially altering the distribution of global capacity factors and capital recovery factors, thereby influencing future global average LCOE forecasts (Polzin et al., 2021; Lonergan et al., 2023; Egli et al., 2019).

#### IV.6.7 IAM scenario comparison

We apply our forecasts for the total investment costs  $\mathcal{I}_j$  to different IAM scenarios from the AR6 IPCC report database (Byers et al., 2022). The purpose of this exercise is to compare the forecasts for future national investment costs to the costs projected by respective IAM scenarios and identify potential issues with the latter.

Our forecasts for  $\mathcal{I}_j$  are conditional on global solar and wind deployment. Since we require annual deployment scenario data, we consider the global solar and wind deployment in 2040 and interpolate linearly with the global solar and wind deployment in 2020, our base year. The reason for this approach over using the 5-year scenario data directly is that a majority of scenarios show a capacity in 2020 that is below what was actually observed and total investment costs that are larger than today's costs. If we were to use the scenario data as it is, it would not be a fair comparison since our forecasts are sensitive to these input values. Next to linear interpolation, we can also interpolate the log-deployment. We omit from showing the results of this analysis here since there is no material difference from the linear interpolation.

For each possible value of future renewable deployment, we forecast the associated

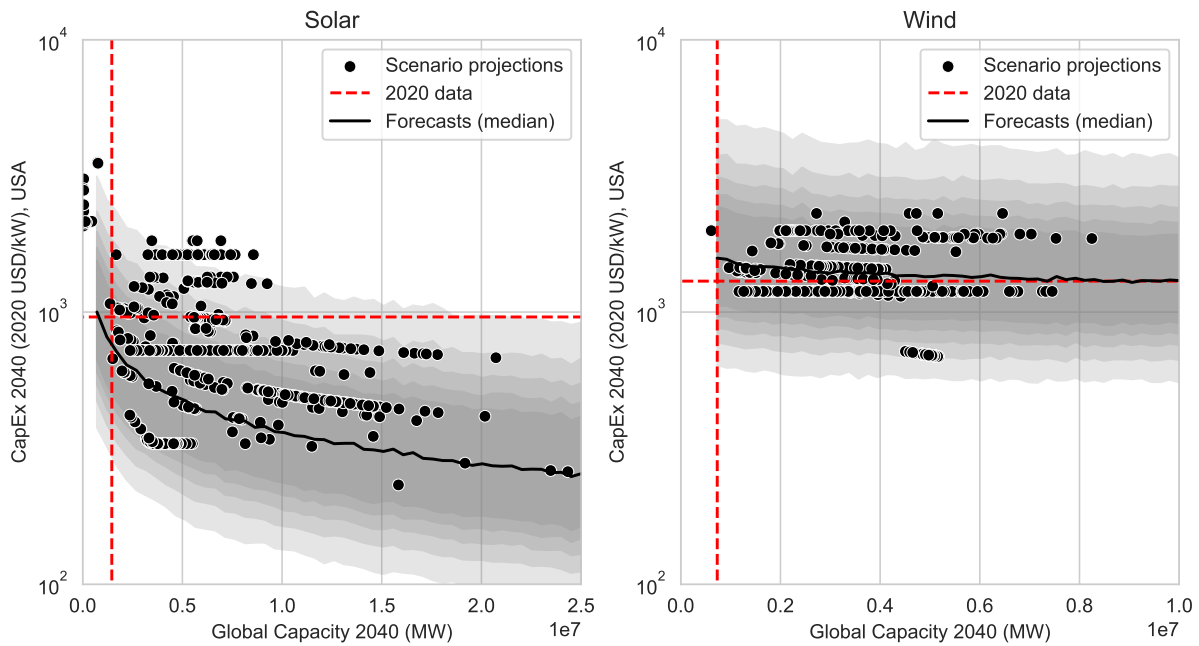


Figure IV.5: **Solar and wind total investment cost projections reported in IAM scenarios against probabilistic projections from our model.** The dots show the total investment cost projections for the US in different IAM scenarios. The x-axis shows the associated global solar / wind capacity. We compare them with the investment cost projections obtained from our model. The gray areas show 50-90% confidence intervals in 10% steps. The blue line shows the median cost from our model. The red dotted lines show the global capacity and investment cost in the US for 2020. For wind, the IAM projections generally align well with the median forecast. For solar, over 80% of projected IAM costs are above the median costs forecast through our model. In both cases, the probabilistic forecasts show significantly more uncertainty around the median forecasts than the IAM projections represent.

total investment costs and compare them to the IAM scenario projections. Figure IV.5 shows our forecasts and the scenario projections, conditional on the global solar and wind deployment. For wind, we see that the IAM projections align well with our forecasts. For solar, there is less alignment, with 80% of scenario costs above our median forecast. The reason for this is that future solar costs are sensitive to their current costs. IAM scenarios are quickly outdated, leading to these discrepancies. In some cases, the scenarios even project renewable capacity in 2040 to be lower than today. In both cases, our forecasts show substantially more uncertainty than the set of IAM scenarios. For solar, over 90% of scenario costs are above our 25th forecast percentile, for wind it is over 80% of scenarios.

## References

- Barbose, Galen, Naïm Richard Darghouth, Eric O’Shaughnessy, and Forrester**, “Tracking the Sun 2024 Edition,” Technical Report, Lawrence Berkley National Laboratory August 2024.
- Baumgärtner, C. Lennart, Rupert Way, Matthew C. Ives, and J. Doyne Farmer**, “The Need for Better Statistical Testing in Data-Driven Energy Technology Modeling,” *Joule*, September 2024, 8 (9), 2453–2466.
- Boccard, Nicolas**, “Capacity Factor of Wind Power Realized Values vs. Estimates,” *Energy Policy*, July 2009, 37 (7), 2679–2688.
- Bolinger, Mark and Joachim Seel**, “Utility-Scale Solar: Empirical Trends in Project Technology, Cost, Performance, and PPA Pricing in the United States – 2018 Edition,” Technical Report 1477381 September 2018.
- , – , **Cody Warner, and Dana Robson**, “Utility-Scale Solar, 2022 Edition: Empirical Trends in Deployment, Technology, Cost, Performance, PPA Pricing, and Value in the United States,” Technical Report None, 1888246, ark:/13030/qt7496x1pc September 2022.
- , **Ryan Wiser, and Eric O’Shaughnessy**, “Levelized Cost-Based Learning Analysis of Utility-Scale Wind and Solar in the United States,” *iScience*, June 2022, 25 (6), 104378.
- Boretti, Alberto and Stefania Castelletto**, “Trends in Performance Factors of Wind Energy Facilities,” *SN Applied Sciences*, September 2020, 2 (10), 1718.
- Byers, Edward, Volker Krey, Elmar Kriegler, Keywan Riahi, Roberto Schaeffer, Jarmo Kikstra, Robin Lamboll, Zebedee Nicholls, Marit Sandstad, Chris Smith, van der Wijst, Kaj, Al -Khourdajie, Alaa, Franck Lecocq, Portugal-Pereira, Joana, Yamina Saheb, Anders Stromman, Harald Winkler, Cornelia**

Auer, Elina Brutschin, Matthew Gidden, Philip Hackstock, Mathijs Harm-  
sen, Daniel Huppmann, Peter Kolp, Claire Lepault, Jared Lewis, Giacomo  
Marangoni, Müller-Casseres, Eduardo, Ragnhild Skeie, Michaela Werning,  
Katherine Calvin, Piers Forster, Celine Guivarch, Tomoko Hasegawa, Malte  
Meinshausen, Glen Peters, Joeri Rogelj, Bjorn Samset, Julia Steinberger,  
Massimo Tavoni, and van Vuuren, Detlef, “AR6 Scenarios Database,” November  
2022.

Creutzig, Felix, Peter Agoston, Jan Christoph Goldschmidt, Gunnar Luderer,  
Gregory Nemet, and Robert C. Pietzcker, “The Underestimated Potential of Solar  
Energy to Mitigate Climate Change,” *Nature Energy*, August 2017, 2 (9).

Egli, Florian, Bjarne Steffen, and Tobias S. Schmidt, “A Dynamic Analysis of  
Financing Conditions for Renewable Energy Technologies,” *Nature Energy*, December  
2018, 3 (12), 1084–1092.

– , – , and – , “Bias in Energy System Models with Uniform Cost of Capital Assumption,”  
*Nature Communications*, October 2019, 10 (1), 4588.

Elia, A., M. Taylor, B. Ó Gallachóir, and F. Rogan, “Wind Turbine Cost Reduction:  
A Detailed Bottom-up Analysis of Innovation Drivers,” *Energy Policy*, December 2020,  
147, 111912.

Elshurafa, Amro M., Shahad R. Albardi, Simona Bigerna, and Carlo Andrea  
Bollino, “Estimating the Learning Curve of Solar PV Balance-of-System for over 20  
Countries: Implications and Policy Recommendations,” *Journal of Cleaner Production*,  
September 2018, 196, 122–134.

Farmer, J. Doyne and Arjun Makhijani, “A US Nuclear Future? Not Wanted, Not  
Needed,” *Nature*, September 2010, 467 (7314), 391–393.

– and François Lafond, “How Predictable Is Technological Progress?,” *Research Policy*,  
April 2016, 45 (3), 647–665.

- Funk, Jeffrey L. and Christopher L. Magee**, “Rapid Improvements with No Commercial Production: How Do the Improvements Occur?,” *Research Policy*, April 2015, 44 (3), 777–788.
- Guivarch, Céline, Thomas Le Gallic, Nico Bauer, Panagiotis Fragkos, Daniel Huppmann, Marc Jaxa-Rozen, Ilkka Keppo, Elmar Kriegler, Tamás Krisztin, Giacomo Marangoni, Steve Pye, Keywan Riahi, Roberto Schaeffer, Massimo Tavoni, Evelina Trutnevyte, Detlef van Vuuren, and Fabian Wagner**, “Using Large Ensembles of Climate Change Mitigation Scenarios for Robust Insights,” *Nature Climate Change*, May 2022, 12 (5), 428–435.
- Helveston, John Paul, Gang He, and Michael R. Davidson**, “Quantifying the Cost Savings of Global Solar Photovoltaic Supply Chains,” *Nature*, December 2022, 612 (7938), 83–87.
- Huenteler, Joern, Christian Niebuhr, and Tobias S. Schmidt**, “The Effect of Local and Global Learning on the Cost of Renewable Energy in Developing Countries,” *Journal of Cleaner Production*, August 2016, 128, 6–21.
- IRENA**, “The Cost of Financing for Renewable Power,” Technical Report, International Renewable Energy Agency, Abu Dhabi 2023.
- IRENA**, “Renewable Power Generation Costs in 2023,” Technical Report, IRENA, Abu Dhabi 2024.
- Key, Alicia, Owen Robers, and Annika Eberle**, “Scaling Trends for Balance-of-System Costs at Land-Based Wind Power Plants: Opportunities for Innovations in Foundation and Erection,” *Wind Engineering*, 2022, 46 (3), 896–913.
- Komendantova, Nadejda, Thomas Schinko, and Anthony Patt**, “De-Risking Policies as a Substantial Determinant of Climate Change Mitigation Costs in Developing Countries: Case Study of the Middle East and North African Region,” *Energy Policy*, April 2019, 127, 404–411.

**Kothari, Sumit, Nadia Ameli, and Michael Grubb**, “Drivers of Balance-of-System Costs for Solar PV Technology,” 2023.

**Lacal-Aránzategui, Roberto, José M. Yusta, and José Antonio Domínguez-Navarro**, “Offshore wind installation: Analysing the evidence behind improvements in installation time,” *Renewable and Sustainable Energy Reviews*, September 2018, *92*, 133–145.

**Lafond, François**, “Forecasting Technological Progress,” March 2025.

– , **Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J. Doyne Farmer**, “How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts,” *Technological Forecasting and Social Change*, March 2018, *128*, 104–117.

– , **Diana Seave Greenwald, and J. Doyne Farmer**, “Can Stimulating Demand Drive Costs Down? World War II as a Natural Experiment,” *SSRN Electronic Journal*, June 2020.

**Lonergan, Katherine E., Florian Egli, Sebastian Osorio, Giovanni Sansavini, Michael Pahle, Tobias S. Schmidt, and Bjarne Steffen**, “Improving the Representation of Cost of Capital in Energy System Models,” *Joule*, March 2023, *7* (3), 469–483.

**Lutz, Christian, Markus Flaute, Ulrike Lehr, and Kirsten S. Wiebe**, “Economic Impacts of Renewable Power Generation Technologies and the Role of Endogenous Technological Change,” *Open Journal of Applied Sciences*, November 2015, *5* (11), 696–704.

**Masson, Gaetan and Izumi Kaizuka**, “Trends in Photovoltaic Applications 2022,” Technical Report 43:2022, IEA PVPS TCP 2022.

**McNerney, James, J. Doyne Farmer, and Jessika E. Trancik**, “Historical Costs of Coal-Fired Electricity and Implications for the Future,” *Energy Policy*, June 2011, *39* (6), 3042–3054.

**Meng, Jing, Rupert Way, Elena Verdolini, and Laura Diaz Anadon**, “Comparing Expert Elicitation and Model-Based Probabilistic Technology Cost Forecasts for the Energy Transition,” *Proceedings of the National Academy of Sciences of the United States of America*, July 2021, *118* (27).

**Mulugetta, Yacob, Youba Sokona, Philipp A. Trotter, Samuel Fankhauser, Jessica Omukuti, Lucas Somavilla Croxatto, Bjarne Steffen, Meron Tesfamichael, Edo Abraham, Jean-Paul Adam, Lawrence Agbemabiese, Churchill Agutu, Mekalia Paulos Aklilu, Olakunle Alao, Bothwell Batidzirai, Getachew Bekele, Anteneh G. Dagnachew, Ogunlade Davidson, Fatima Denton, E. Ogheneruona Diemuodeke, Florian Egli, Eshetu Gebrekidan Gebresilassie, Muluaem Gebresilassie, Mamadou Goundiam, Haruna Kachalla Gujba, Yohannes Hailu, Adam D. Hawkes, Stephanie Hirmer, Helen Hoka, Mark Howells, Abdulrasheed Isah, Daniel Kammen, Francis Kemausuor, Ismail Khennas, Wikus Kruger, Ifeoma Malo, Linus Mofor, Minette Nago, Destenie Nock, Chukwumerije Okereke, S. Nadia Ouedraogo, Benedict Probst, Maria Schmidt, Tobias S. Schmidt, Carlos Shenga, Mohamed Sokona, Jan Christoph Steckel, Sebastian Sterl, Bernard Tembo, Julia Tomei, Peter Twesigye, Jim Watson, Harald Winkler, and Abdulmutalib Yussuff**, “Africa Needs Context-Relevant Evidence to Shape Its Clean Energy Future,” *Nature Energy*, October 2022, pp. 1–8.

**Nagy, Bela, J. Doyne Farmer, Quan M. Bui, and Jessika E. Trancik**, “Statistical Basis for Predicting Technological Progress,” *PLoS ONE*, February 2013, *8* (2), 52669.

**Nemet, Gregory F.**, “Beyond the Learning Curve: Factors Influencing Cost Reductions in Photovoltaics,” *Energy Policy*, November 2006, *34* (17), 3218–3232.

- , **Jiaqi Lu, Varun Rai, and Rohan Rao**, “Knowledge Spillovers between PV Installers Can Reduce the Cost of Installing Solar PV,” *Energy Policy*, September 2020, *144*, 111600.
- Peters, Glen P., Alaa Al Khourdjie, Ida Sognaes, and Benjamin M. Sanderson**, “AR6 Scenarios Database: An Assessment of Current Practices and Future Recommendations,” *npj Climate Action*, September 2023, *2* (1), 1–6.
- Polzin, Friedemann, Mark Sanders, Bjarne Steffen, Florian Egli, Tobias S. Schmidt, Panagiotis Karkatsoulis, Panagiotis Fragkos, and Leonidas Parousos**, “The Effect of Differentiating Costs of Capital by Country and Technology on the European Energy Transition,” *Climatic Change*, July 2021, *167* (1-2).
- Raupach-Sumiya, Jörg, Hironao Matsubara, Andreas Prah, Astrid Aretz, and Steven Salecki**, “Regional Economic Effects of Renewable Energies - Comparing Germany and Japan,” *Energy, Sustainability and Society*, December 2015, *5* (1), 10.
- Riva, Alberto Dalla, Janos Hethey, Silke Luers, Anna-Kathrin Wallasch, Knud Rehfeldt, Aidan Duffy, David Edward Weir, Maria Stenkvis, Andreas Uihlein, Tyler Stehly, and Eric Lantz**, “IEA Wind TCP Task 26: Wind Technology, Cost, and Performance Trends in Denmark, Germany, Ireland, Norway, Sweden, the European Union, and the United States: 2008-2016,” Technical Report, National Renewable Energy Laboratory (NREL), Golden, CO (US) November 2018.
- Steffen, Bjarne, Martin Beuse, Paul Tautorat, and Tobias S. Schmidt**, “Experience Curves for Operations and Maintenance Costs of Renewable Energy Technologies,” *Joule*, February 2020, *4* (2), 359–375.
- Wang, Rong, Sandra Hasanefendic, Elizabeth Von Hauff, and Bart Bossink**, “The Cost of Photovoltaics: Re-evaluating Grid Parity for PV Systems in China,” *Renewable Energy*, July 2022, *194*, 469–481.

**Way, Rupert, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer**, “Empirically Grounded Technology Forecasts and the Energy Transition,” *Joule*, September 2022, *6*, 1–26.

**Wiser, Ryan, Mark Bolinger, and Eric Lantz**, “Assessing Wind Power Operating Costs in the United States: Results from a Survey of Wind Industry Experts,” *Renewable Energy Focus*, September 2019, *30*, 46–57.

# Chapter V

## Innovation Bandits: A Dynamic Portfolio Strategy with Endogenous Rewards

C. LENNART BAUMGÄRTNER, LAURIN KÖHLER-SCHINDLER,  
JACQUELYN PLESS

When allocating resources across innovation efforts, managers, investors, policymakers, and system planners face a trade-off between incurring high current costs and achieving future cost reductions, as technologies evolve endogenously with experience and economic conditions. Future innovation is also uncertain, giving rise to a non-linear dynamic portfolio optimization problem. Yet optimization methods typically treat innovation as exogenous and discard endogenous uncertainties. As a result, these methods can be difficult to interpret and overly restrictive in their solution space. In this paper, we develop an alternative dynamic portfolio optimization framework that integrates endogenous learning, showing explicitly how Multi-Armed Bandit (MAB) strategies can be applied to technology portfolios with empirically validated experience curves. MAB-based approaches explicitly capture the exploration-exploitation trade-off, dynamically adjusting investment decisions in response to observed cost declines. In our setting, we implement a MAB strategy based on the Gittins index. The resulting adaptive strategy prioritises investments with volatile future costs over those with less volatility. We demonstrate our method through an application to the U.S. light vehicle sector, illustrating how dynamic MAB-based strategies can outperform static alternatives that ignore endogenous uncertainties under a set of assumptions. Taken together, our findings suggest that adaptive strategies may facilitate more rapid market transitions, providing important implications for both innovation policy design and private investment.

---

**Acknowledgments:** For helpful comments and feedback, we thank J. Doyne Farmer, Rupert Way, Cameron Hepburn, Alex Teytelboym, and seminar and conference participants at the EPOC conference 2024. We are grateful to Tiziano De Angelis, Giorgio Ferrari, Alessandro Milazzo, Neofytos Rodosthenous, Mateo Rodríguez Polo and members of Group 3 at the Department of Mathematics, ETH Zurich, for helpful comments and feedback on the analytic approach to calculate the Gittins index. We also thank the Institute of New Economic Thinking, Oxford, for their generous financial support. L. KS. has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant 851565).

## V.1 Introduction

Innovation is a key driver of firm performance and competitiveness (Chao and Kavadias, 2008) as well as aggregate economic growth (Solow, 1957; Romer, 1990). In practice, fostering innovation requires decision-makers—managers, investors, policymakers, and system planners—to make complex decisions about how to allocate resources across research and development (R&D) projects, technologies, and policy initiatives. Doing so also entails considering allocations over time and developing dynamic investment plans that are continuously reassessed and redirected as technologies and markets evolve, requiring decision-makers to balance short-term costs with uncertain longer-term returns. This challenge is particularly acute for nascent technologies, which play a crucial role in steering the *direction* of innovation toward addressing pressing global challenges such as climate change.

However, many common portfolio optimization and system planning strategies fall short on capturing the intrinsic nature of innovation and the practicalities of managing innovation portfolios, leading to sub-optimal investment strategies. In particular, many models assume exogenous technological change. But a key feature of technological progress is that costs tend to decline endogenously with cumulative investment and deployment—a relationship formalized by experience curves (Wright, 1936; Nagy et al., 2013; Lafond et al., 2018). This creates an inherent trade-off: investing in emerging technologies may be costly today but accelerate future cost reductions. Moreover, future cost reductions can be uncertain.<sup>1</sup> Decision-makers must therefore balance *exploration* (investing in riskier, less mature technologies with high learning potential) against *exploitation* (focusing on established technologies with more predictable returns).

Traditional portfolio optimization techniques do not adequately capture these interdependencies. Most models assume static decision-making and exogenous technological change. When innovation is endogenous, the optimization problem becomes non-convex

---

<sup>1</sup>As we will see later, uncertain does not imply unpredictable. Within a range of uncertainty, we can make specific predictions of future cost reductions (Farmer and Lafond, 2016).

(Way et al., 2019), producing solutions that are non-unique and difficult to interpret (Farmer et al., 2015). While the use of experience curves is increasingly common in static technology portfolios (Krey, 2014), optimizing *dynamic* resource allocation remains challenging due to the high dimensionality of the solution space (Powell, 2011). These approaches are computationally intensive, thus limiting the temporal resolution and granularity of portfolio choices represented. Moreover, static approaches typically assume that innovation follows a deterministic path rather than a stochastic process, further limiting their real-world applicability.

In this paper, we introduce a new approach to dynamic innovation portfolio optimization that overcomes some of these limitations. We develop a framework that incorporates endogenous learning based on the Multi-Armed Bandit (MAB) problem, a well-established tool for addressing dynamic exploration-exploitation trade-offs (Gittins et al., 2011). Although MAB-based strategies have been widely applied in resource allocation and reinforcement learning problems, their application to technology investment with realistic innovation dynamics remains under-explored. We follow one specific MAB solution to define novel portfolio strategies for technologies that follow experience curves, resulting in a computationally efficient solution to the problem of allocating resources across competing technologies. While our choice of MAB solution comes with a number of practical restrictions, it is the first explicit application of the MAB framework to innovation investments. As such, results also could be applied to other innovation drivers (e.g., R&D expenditures) if practical limitations to the presented approach are overcome. Our framework is thus a first step towards a broader practical toolkit for managers, investors, system planners, and policymakers seeking to optimize innovation investments. To illustrate its practical relevance, we apply it to the light passenger vehicles context to determine the optimal tipping point strategy for transitioning away from fossil fuel-based passenger transportation.

At the core of our approach is one of the most well-known phenomena underlying innovation dynamics, Wright’s law (Wright, 1936), whereby technology costs improve with

cumulative deployment. Recent empirical studies of experience curves describe innovation as a stochastic process akin to a semimartingale. The resulting model captures long-term experience effects, stochastic uncertainty in future cost reductions, and model uncertainty associated with each technology. [Lafond et al. \(2018\)](#) show that the stochastic experience curve holds empirically in their study of around 60 technologies. Furthermore, [Meng et al. \(2021\)](#) show that the resulting predictions outperform other approaches to forecasting future technology costs with respect to long-term accuracy, particularly relative to expert elicitations. See [Lafond \(2025\)](#) for a comprehensive overview of current developments and future research on this topic.

Experience curves are widely used in economic forecasts and techno-economic models. They especially play an important role in energy system models and integrated assessment models that are used to assess the impact of climate change on the economy ([Krey, 2014](#)), both of which frequently inform energy investment planning and sustainable development policy decisions. In these models, investors or policymakers repeatedly allocate resources to a portfolio of competing technologies to maximize their total expected utility. Prior to every investment decision, information about potential choices changes, such that the allocation is revised dynamically. Decision-makers thus must assess the trade-off between current costs and future cost reductions that may occur through innovation. That is, these are dynamic exploration-exploitation problems.

Given the high dimensionality of the solution space when solving for optimal dynamic resource strategies ([Powell, 2011](#)), techno-economic models typically make two simplifying assumptions. First, they assume that technological innovation is independent of investment choices, and second, that portfolio decisions are static. Under these assumptions, the optimization problem reduces to maximizing expected discounted utility. This is analytically attractive, and both of these assumptions are reasonable when portfolio dynamics are marginal or the impact of an investor's decisions on the underlying portfolio is negligible. However, when considering a portfolio of technologies governed by experience curves, these assumptions break down, resulting in sub-optimal investment strategies.

We take an alternative approach that builds on a rigorous reformulation of the dynamic innovation portfolio problem using MAB theory. By representing each technology’s cost evolution as an Itô-stochastic process (Ito, 1951)—with a time-inhomogeneous drift determined by endogenous learning via experience curves—we recast the classical portfolio optimization challenge into a dynamic exploration-exploitation trade-off. There are multiple solutions to this problem, most prominently index strategies and non-parametric strategies. Here, we determine the optimal exploration-exploitation decision rule through an explicit expression for the Gittins index, which serves as a “fair option price” for independently investing in each technology. This derivation, grounded in established results from stochastic control theory, optimal stopping, and earlier MAB literature (e.g., Mandelbaum (1987), Karatzas (1984), EL Karoui and Karatzas (1994), and Peskir and Shiryaev (2006)), allows us to overcome the “curse of dimensionality” inherent in dynamic models and to solve for optimal investment strategies under a set of assumptions alternative to traditional approaches.

Our numerical solution for the Gittins index exploits a backward-integration scheme to approximate the free boundary that characterizes the underlying optimal stopping problem. This free-boundary formulation not only provides a clear economic interpretation—quantifying the trade-off between immediate costs and potential future cost reductions—but also yields substantial computational advantages. Specifically, the decoupling of the index calculation across technologies implies that the computational effort scales linearly with the number of portfolio choices. This computational efficiency makes our method applicable to a range of practical problems in technology investment and policy design. At the same time, a number of strong assumptions are required to compute the Gittins index that may not hold in practice. In particular, the Gittins index’ performance is contingent on the absence of portfolio switching costs, instantaneous innovation feedback, and full knowledge of experience curve parameters.

Under these assumptions, we find that our model outperforms commonly employed strategies based on the expected net present cost (ENPC) of each technology. Unlike

static approaches, the dynamic portfolio adjusts over time, allowing investors to take on more short-term volatility. Investors can consider their future option to switch resource allocations, making volatile technologies more attractive. If the outcomes of volatile technologies are unfavorable, investors adjust their portfolios. Our Gittins index strategy thus generalizes the well-known result of real options theory that optimal dynamic strategies are risk-seeking rather than risk-neutral (Dixit and Pindyck, 1994, Chapter 5).

To illustrate the practical implications of our strategy, we apply it to an empirical case: capital investment in battery-electric vehicles (BEVs) versus internal combustion engine (ICE) vehicles in the light passenger vehicles context. Decarbonizing transportation is essential to reducing the environmental impact of economic activity, and BEVs are poised to play a central role in this transition (Hoehne et al., 2023). However, upfront capital requirements for BEV remain higher than conventional alternatives (BNEF, 2023). With continued innovation, capital costs could decline, but the optimal investment strategies to facilitate this transition remain an open question.

We apply our strategy to portfolios of technologies evolving under Wright’s law using an empirically validated stochastic learning curve model. We use data on light passenger vehicles in the U.S. to calibrate the model, drawing on historical trends in battery costs and ICE prices, and estimate technology-specific parameters (e.g., learning, noise, and mean reversion). Sufficient historical data on BEVs allows us to avoid large parameter uncertainty. These calibrations inform our computation of the Gittins index, allowing us to simulate dynamic investment decisions that adapt to cumulative deployment and stochastic cost fluctuations. Specifically, it allows us to identify the theoretically optimal tipping point, i.e., the maximum cost of BEVs below which it is optimal to invest in them over ICEs, depending on the learning potential and discount rate.

Our dynamic strategy yields insights that complement traditional static approaches. Specifically, our simulations demonstrate that, under low discount rates, the Gittins index strategy already favors investments in BEVs owing to their potential for rapid cost reductions through cumulative learning and endogenous volatility. In contrast, static

ENPC strategies remain locked into ICE investments, missing the opportunity to exploit emerging cost declines in BEVs. Overall, our results indicate that the dynamic strategy not only achieves a lower mean net present cost but also provides a flexible response mechanism to technological and market uncertainties, thereby offering policy-relevant insights for accelerating the transition toward cleaner transportation technologies.

### V.1.1 Related Portfolio Optimization Literature

Given the ubiquity of resource allocation problems across various settings, the portfolio optimization literature is unsurprisingly vast. We focus our attention here on methods commonly applied in the innovation context.

Our methods are rooted in mathematical finance and stochastic control theory. If we were to assume that technological change is exogenous and portfolio decisions are static, traditional Markowitz portfolio theory would be appropriate for optimizing the (typically convex) mean-variance trade-off, as long as the utility function is approximately quadratic and the noise is approximately normal (Markowitz, 1952, 2014). The associated methods can be generalized to dynamic settings in discrete (Samuelson, 1969) and continuous time (Merton, 1969). However, when relaxing the exogeneity assumption to include endogenous rewards (i.e., innovation), the Markowitz optimization problem becomes non-convex (Way et al., 2019), leading to solutions that are no longer unique or robust and difficult to interpret (Farmer et al., 2015).

A variety of solutions to this challenge have been developed that are specific to different settings. Examples include myopic models that deal with temporal uncertainty through limited foresight (e.g., Keppo and Strubegger (2010)), stochastic models that deal with parameter uncertainty in a Monte Carlo fashion (e.g., Messner et al. (1996); Milford et al. (2022)), or approximate dynamic programming models applying forward (e.g., Kumbaroğlu et al. (2008)) and backward (e.g., Bukenberger and Palmintier (2018)) induction methods.<sup>2</sup>

However, solutions so far come with a different set of limitations. Two models that

---

<sup>2</sup>See also Ioannou et al. (2017) for a review of stochastic methods in energy planning models.

are particularly relevant for our work are those in [Way et al. \(2019\)](#), who explore the full solution space of static technology portfolios subject to experience curves, and [Cowan \(1991\)](#), who introduces MAB concepts for technology portfolios but does not apply an explicit optimization strategy. So far, these non-linear models have been limited by the stochastic processes they can represent and the size of the dynamic solution space. Additionally, they are computationally intensive, thus limiting the temporal resolution and granularity of portfolio choices represented.

To the best of our knowledge, no models are employing MAB concepts to study technology portfolios while also representing innovation dynamics as a stochastic process. Yet, empirical data shows that technology performance measures (such as costs and efficiency) subject to innovation tend to behave like semimartingales.<sup>3</sup> This intuition has been shown empirically using learning curves and experience curves ([Farmer and Lafond, 2016](#); [Lafond et al., 2018](#); [Lafond, 2025](#)), whereby the cost of a technology decreases with time or cumulative deployment, respectively. One possible mechanism giving rise to this universal power-law behavior is a random-search model over technologies with different degrees of complexity ([McNerney et al., 2011b](#)). For this work, we highlight applications for portfolios evolving under experience curves, but similar methods can be applied to other innovation processes, such as economies of scale, R&D investments, or combinations thereof ([Meng et al., 2021](#); [Nagy et al., 2013](#)).

MAB-based strategies have been applied to resource allocation problems since the emergence of dynamic programming in the 1950s ([Gittins, 1979](#); [Gittins et al., 2011](#)). Lately, the success of machine learning, especially reinforcement learning, has led to a resurgence of the method in both computer science and stochastic programming. There are two prominent strands of the MAB literature: *index strategies* and *non-parametric strategies* ([Scott, 2010](#)).

Index strategies provide explicitly optimal dynamic choices when the underlying

---

<sup>3</sup>A semimartingale is the sum of a local martingale and a càdlàg adapted process of locally bounded variation. We expect innovation processes to be subject to a non-trivial drift that is predictable based on past data. Martingales are, therefore, too restrictive to describe these processes. Here, we use Itô-stochastic processes to describe innovation.

stochastic processes are known. The most prominent index strategy is the Gittins index (Gittins, 1979), although important alternatives, such as the Whittle index (Whittle, 1988), exist for other settings. In management science, Gittins index strategies have been applied to risk-averse portfolio choices (Malekipirbazari and Çavuş, 2024) and stochastic optimization with distrust in the underlying model (Kim and Lim, 2016). In this setting, the exploration-exploitation trade-off refers to an investor exploring portfolio choices with high expected future rewards under known uncertainty versus exploiting portfolio choices with high current rewards. When exploring portfolio choices, the investor updates their uncertainty about future rewards through gained *experience*.

Non-parametric strategies, particularly Thompson-sampling, upper-confidence bounds,  $\epsilon$ -greedy strategies, and variations thereof, are asymptotically optimal strategies that are both simple to use and accurate in settings where the underlying processes are unknown (Agrawal and Goyal, 2012). They are widely used in reinforcement learning but also in operations research, such as in Cheung et al. (2022)'s investigation of non-stationary bandits with bounded drift. In this context, the exploration-exploitation trade-off refers to an investor exploring portfolio choices with unknown rewards versus exploiting portfolio choices with known expected rewards. When exploring portfolio choices, the investor updates their expectation or prior of future rewards through *learning* without influencing the underlying dynamics.

In this paper, we rely on Gittins index strategies. Given the universality of experience curves, we consider the stochastic process describing a technology's cost evolution to be known; technology-specific model parameters can be estimated empirically. Furthermore, we take additional assumptions, particularly the absence of switching costs and immediate feedback in a continuous time setting. This allows us to use existing MAB results (e.g., EL Karoui and Karatzas (1994), Mandelbaum (1987), Karatzas (1984) and others) to find explicit Gittins index strategies that are both theoretically optimal and computationally efficient. Given the strong assumptions required, there are a number of avenues to generalise our results based on the Gittins index or apply alternative approaches based on

non-parametric MAB solutions.

## V.1.2 Summary of Contributions

Our paper makes five key contributions. Firstly, our methodological advances bridge gaps between the mathematical finance, innovation management, and technological change literatures by incorporating endogenous learning into portfolio optimization. Innovation is an inherently endogenous and uncertain process but modeling endogenous uncertainties has traditionally proven challenging in portfolio optimization (Salo et al., 2024). For contexts where cumulative investments drive cost reductions, we leverage empirically validated stochastic models of technology cost evolution (e.g., Lafond et al. (2018) and others based on Wright’s law) and develop a framework that delivers dynamic investment strategies that outperform conventional techniques, such as static strategies based solely on ENPC, under a set of assumptions. Our approach can be applied to other drivers of innovation beyond cumulative investment in future research.

Secondly, we extend existing multi-armed bandit (MAB) results by deriving an explicit expression for the Gittins index for a class of time-inhomogeneous stochastic processes with monotone drift. Our derivation synthesizes foundational work from the early MAB literature (e.g., Mandelbaum (1987); Karatzas (1984); Kaspi and Mandelbaum (1995), EL Karoui and Karatzas (1994)) with more recent advances in stochastic control theory and optimal stopping (e.g., Peskir and Shiryaev (2006); Milazzo (2024); Peskir (2005); de Angelis and Peskir (2020)). The corresponding numerical implementation—based on a backward-integration scheme—yields an asymptotically optimal and computationally efficient strategy that scales linearly with the number of portfolio choices. This paves the way for future extensions—such as accounting for parameter uncertainty and portfolio constraints—thus offering a versatile tool for both academic research and practical decision-making in high stakes, uncertain environments.

Thirdly, our method connects emerging developments in machine learning and artificial intelligence, real options theory in investment decisions, innovation and technological

change, and the broader portfolio optimization literatures. For instance, recent work by [Gunjan and Bhattacharyya \(2023\)](#) highlights the role of reinforcement learning in dynamic portfolio strategies. From a MAB perspective, this topic is closely related to contextual learning ([Bouneffouf et al., 2020](#)) and Multi-Armed Bandits with dependent arms ([Shen et al., 2015](#); [Huo and Fu, 2017](#); [Malekipirbazari and Çavuş, 2024](#); [Sunar and Wang, 2024](#)), highlighting how MAB frameworks can be utilized in financial markets with exogenous volatility and risk aversion. Real options theory has been applied to a number of single investment decision questions (e.g., [Décamps et al. \(2005\)](#); [Klein \(2009\)](#); [Grenadier and Malenko \(2010\)](#)). Applying the MAB literature extends this to repeated investment decisions in innovation portfolios. Unlike models that treat innovation deterministically or rely on static optimization, our framework explicitly incorporates the stochastic evolution of technology performance and the value of the associated optionality. Similar to a real options approach, we identify the optimal timing of adjusting resource allocations. In a similar way, our work also relates to the economics literature on strategic experimentation and learning (e.g., [Keller and Rady \(1999\)](#); [Bolton and Harris \(1999\)](#); [Moscarini and Smith \(2001\)](#)).

Fourth, our application to the light vehicle sector and energy transitions illustrates the practical relevance of the MAB approach in an important and timely context. Given how transitioning away from fossil fuel-based transportation will be critical for addressing climate change and reducing the environmental impact of economic activity, we provide direct practical insight for ongoing public policy discussions (e.g., [Whitehead et al. \(2023\)](#); [IEA \(2024\)](#); [Milovanoff et al. \(2020\)](#)). We reveal how adaptive strategies can identify critical tipping points that can trigger market transformations ([Geels and Ayoub, 2023](#)). Our results may be especially of interest to the energy system modeling community (e.g., [Unterluggauer et al. \(2022\)](#); [Hoehne et al. \(2023\)](#); [Muratori et al. \(2020\)](#)), although our methods can be readily applied to other settings where there are highly substitutable but uncertain technology options. At the same time, this example also highlights key limitations with the chosen approach based on the Gittins index. Several strong assumptions are

required in this setting that may not be appropriate in different practical settings.

Finally, from a public planning and policy perspective, our work is pertinent for policymakers and government agencies aiming to design effective innovation policies and support programs. We illustrate how dynamic investment strategies can accelerate cost reductions and market transitions, outperforming static alternatives. Our findings suggest that, with novel technologies, smaller project sizes, fast feedback, and low switching costs, portfolios benefit from learning effects and should be highly concentrated to identify “winners” quickly. This stands in contrast to “neutral” policy choices that do not discriminate between different portfolio options (e.g., [Fabra and Montero \(2023\)](#); [de Mello Santana \(2016\)](#); [Rodrik \(2014\)](#)). Adaptive investment strategies that target tipping points are also highly relevant for policymakers and serve as tools for designing adaptive, cost-effective innovation policies, since sensitive intervention points can trigger market transformations ([Mealy et al., 2023](#); [Farmer et al., 2015](#)). With this in mind, our paper may be of interest and timely for both policymakers and private sector investors interested in socially responsible investment (e.g., see [Ballesterio et al., 2012](#)) amongst others).

This work has multiple caveats that we discuss throughout the paper. Most importantly, the Gittins strategy required multiple strong assumptions, such as known learning parameters, immediate feedback, no switching costs and substitutable technologies. Nevertheless, the framework presents a novel alternative to current approaches to this problem and opens new questions for future research.

### V.1.3 Outline

We start by formulating the optimization problem in Section [V.2](#). We then introduce MAB-based strategies in Section [V.3](#) and show how index strategies can be computed for technology portfolios under endogenous learning. In Section [V.4](#), we apply our strategy to light passenger vehicle investments before discussing the implications and future extensions of this method in Section [V.5](#).

## V.2 A portfolio optimization problem with endogenous rewards

We aim to solve a dynamic allocation problem, whereby future investments are made into different technologies. We choose an objective function that maximizes total utility. The technologies are substitutable and follow independent experience curves with known parameters. We work in continuous time given its analytical advantages.<sup>4</sup> As we will see, this assumption still allows for a meaningful analysis of discrete empirical data. In the main text, we focus on technologies evolving according to Wright’s law. In Appendix C.2.1, we also treat the case of technology costs evolving according to a mean-reverting AR(1) process. Both of these are relevant to our transportation application.

### V.2.1 Technology Cost Evolution

Studies employing MAB-based strategies often make ad-hoc assumptions about the stochastic process that a technology’s cost follows or develop approaches that work well when the noise-generating process is unknown. We take an alternative approach by leveraging an empirically validated stochastic model with endogenous learning and uncertainty that has been shown to reliably predict costs for around 60 technologies. Specifically, Lafond et al. (2018) show that the cost of a technology  $C(t)$  tends to evolve as

$$C(t+1) = C(t) \cdot \left( \frac{Z(t+1)}{Z(t)} \right)^{-\omega} \cdot e^{\eta_t}, \quad (\text{V.1})$$

where  $\omega$  denotes the technology-specific learning parameter,  $Z(t)$  denotes the cumulative deployment of the technology at time  $t \in \mathbb{N}$  in years,  $Z(t)$  increases with the demand in  $(t, t+1]$ ,  $D_{t+1}$ , such that  $Z(t+1) = Z(t) + D_{t+1}$ , and  $\eta_t$  denotes a normally distributed

---

<sup>4</sup>The MAB framework was originally posed in discrete time but later extended to the continuous setting (see Gittins et al. (2011), Mandelbaum (1987), and Karatzas (1984)). Nevertheless, the generalization to discrete portfolio choices is not trivial since we want to allow for a diversified portfolio at each point in time.

noise shock applied in each time period. Taking the logarithm of this expression, we obtain

$$\log C(t+1) = \log C(t) - \omega \cdot \log \left( \frac{Z(t+1)}{Z(t)} \right) + \eta_t. \quad (\text{V.2})$$

Now,  $c(t) \equiv (\log(C(t)))_{t \geq 0}$  is a random walk with  $\mathcal{N}(-\omega \cdot \log \left( \frac{Z(t+1)}{Z(t)} \right), \sigma^2)$  distributed steps.

Empirical data tend to be in discrete time, but it is helpful to work in continuous time to simplify the MAB framework described later. We thus model the logarithm of the technology's cost as a continuous semimartingale  $(c(t))_{t \geq 0}$  defined by

$$c(t) = c_0 - \omega \int_0^t \frac{Z'(s)}{Z(s)} ds + \sigma \cdot B(t) \quad (\text{V.3})$$

$$= c_0 - \omega \log \left( \frac{Z(t)}{Z(0)} \right) + \sigma \cdot B(t), \quad (\text{V.4})$$

where  $(B(t))_{t \geq 0}$  is a standard Brownian motion starting from 0. Equivalently, setting  $\mu(t) := -\omega \frac{Z'(t)}{Z(t)}$ , we can write

$$dc(t) = \mu(t)dt + \sigma dB(t). \quad (\text{V.5})$$

We can recover the original discrete version by observing that

$$c(t+1) = c(t) - \omega \log \left( \frac{Z(t+1)}{Z(t)} \right) + \underbrace{\sigma \cdot (B(t+1) - B(t))}_{\sim \mathcal{N}(0, \sigma^2)}. \quad (\text{V.6})$$

In the following, we make two key assumptions to simplify the calculation of the Gittins index strategy for technologies evolving under this model. Firstly, we will not consider estimation errors in the parameters  $\omega$  and  $\sigma$ . Second, we will assume that the demand  $D$  is constant (i.e., the sum of cumulative deployment over all technologies grows linearly in time). In future work, we hope to relax both assumptions to make the model more widely applicable, see Section [V.5.3](#).

## V.2.2 Objective function

We follow [Mandelbaum \(1987\)](#) and [EL Karoui and Karatzas \(1994\)](#) closely to define the objective function of our dynamic allocation problem. In particular, we choose a utility maximization as the objective, the standard approach when developing index-based MAB solutions. This is not the only option. In particular, non-parametric approaches typically formulate their objective in terms of minimizing the expected *regret* of a strategy, compared to a benchmark. In our case, this could be a static optimization strategy. Future work is required to compare these alternative approaches since they may allow for more realistic assumptions than the setting presented here. The presented utility-based problem formulation fulfils several assumptions required to apply our results for the Gittins index strategy. A full list of the required assumptions is provided in [Appendix C.1](#).

Consider  $d \in \mathbb{N}$  technologies with log-costs  $(c_i(t))_{t \geq 0}$ ,  $1 \leq i \leq d$ , following independent experience curves. Given these choices, we consider dynamic allocation strategies  $\zeta$ . These strategies represent the cumulative production allocated to each technology at each point in time. The allocation depends on currently available information, which includes current and historic costs  $(c_i(t))_{i=1}^d$  and cumulative deployment  $(Z_i(t))_{i=1}^d$ . Mathematically, an allocation strategy is an adapted stochastic process  $\zeta = (\zeta(t))_{t \geq 0}$  taking values in  $\mathbb{R}_+^d$ , whose  $i$ -th coordinate  $\zeta_i(t)$  at time  $t$  models the total production allocated to technology  $i$  within the interval  $[0, t]$ , and which satisfies the following properties:<sup>5</sup>

1. for each  $i = 1, \dots, d$ ,  $\zeta_i = (\zeta_i(t))_{t \geq 0}$  is increasing and starts at 0,
2. for all  $t \geq 0$ ,  $\zeta_1(t) + \dots + \zeta_d(t) = t$ .

Since the strategy represents cumulative production at each point in time, and  $D$  is constant, we can further note that  $Z_i(t) = Z_i(0) + D\zeta_i(t)$ . Here, we also use our assumption that all technologies are perfect substitutes which will be more appropriate in some markets than others (e.g., vehicles or electricity).

---

<sup>5</sup>See Definition 5.1 in [EL Karoui and Karatzas \(1994\)](#) for details.

We aim to find the optimal strategy  $\zeta^*$  achieving the maximum expected utility.<sup>6</sup> Following [Gittins et al. \(2011\)](#), we can write this objective as

$$\sup_{\zeta} \mathbb{E}_{0,c} [U(\zeta)] = \sup_{\zeta} \sum_{i=1}^d \int_0^{\infty} e^{-\alpha t} \mathbb{E}_{0,c} \left[ u_i(c_i(\zeta_i(t))) \dot{\zeta}_i(t) \right] dt \quad (\text{V.7})$$

Here,  $\mathbb{E}_{0,c}$  denotes the expectation under  $c_i(0) = c_i$  for every  $1 \leq i \leq d$ . We write  $\dot{\zeta}(t)$  for the derivative with respect to  $t$ . Since  $\sum_{i=1}^d \dot{\zeta}_i(t) = 1$ , the latter formulation highlights that an allocation strategy corresponds to choosing the shares  $\dot{\zeta}_1(t), \dots, \dot{\zeta}_d(t)$  of the allocation among the  $d$  diffusion processes and collecting rewards according to these shares.  $\alpha > 0$  expresses the discount rate applied to portfolio returns.

The classical index-based MAB framework imposes a utility function that is linear in technology choices and uniformly discounted. For now, we exclude non-linear utility functions, such as log-utility on total consumption or risk-averse utility functions that aim to minimize downside risk on the total portfolio. Instead, we focus on net present costs, where  $u(c(t)) = -C(t) = -e^{c(t)}$  for all technologies. This is consistent with the social welfare function in standard neoclassical growth models (e.g., Ramsey-Cass-Koopmans and successors), where the discount factor represents a time-preference for immediate rewards and which is commonly used in techno-economic models.<sup>7</sup> While exponential discounting is difficult to remove from our setting, this assumption can be relaxed towards bandit-specific discounting or stochastic discounting in the future ([Gittins et al., 2011](#)).

We can also see that the technology costs  $c_i(t)$  depend exclusively on the allocated time  $\zeta_i(t)$ . They evolve iff  $\dot{\zeta}_i(t) > 0$ . This is a slight deviation from [Lafond et al. \(2018\)](#) where the noise shocks  $\eta_t$  are independent of the production of the technology. We discuss further extensions to this in [Section V.5](#).

The two key assumptions in this problem definition are that there are no costs associated to  $\zeta_i$  (or  $\dot{\zeta}_i$ ) and no delay in our knowledge of  $c_i(t)$ . This means that there are no costs associated to switching between technologies, and we receive immediate feedback from an

---

<sup>6</sup>In some settings, maximizing *reward* may be more appropriate phrasing than utility. We use utility to highlight the universality of our setting.

<sup>7</sup>For a more general introduction to utility theory, see for example [Ingersoll \(1987\)](#).

investment. In some cases, these may be appropriate assumptions but in many practical applications they are not. This is a central topic for future research on this topic.

### V.3 Index strategies for continuous innovation bandits

We can use existing MAB methods to find the optimal strategy  $\zeta^*$  in this setting. The key to a successful strategy is to identify the trade-off between exploiting technologies that have low current costs versus exploring technologies that have significant potential for future cost reductions based on their experience curves and associated uncertainty. We can use the *Gittins index* to quantify this trade-off since we know the Itô stochastic utility evolution of the experience curve.

#### V.3.1 The Gittins Index

First proposed by [Gittins \(1979\)](#), the Gittins index theorem states that the optimal investment strategy allocates resources according to an index  $M_i(t)$  assigned to each technology  $i$  at each point in time  $t$  that is independent of all other technologies.

---

**Algorithm 1:** Gittins index strategy

---

For each technology  $1 \leq i \leq d$  at time  $t_0 \geq 0$ , calculate the index  $M_i(t) \in \mathbb{R}$ .

**while**  $t > t_0$  **do**

Find technology  $j$ , with  $M_j(t) = \max_i \{M_i(t)\}$ .

**while**  $M_j(t) = \max_i \{M_i(t)\}$  **do**

Invest in technology  $j$ , i.e.  $\dot{\zeta}_j(t) > 0$

Update the index  $M_j(t)$

---

The existence, uniqueness, and optimality of Gittins index strategies in the continuous setting are established in [Mandelbaum \(1987, Theorem 4\)](#) and [EL Karoui and Karatzas \(1994, Theorem 8.1\)](#)<sup>8</sup>. For the purpose of this work, we follow [EL Karoui and Karatzas](#)

---

<sup>8</sup>Theorem 8.1 establishes optimality of Gittins index strategies in a general, non-Markovian setting. The Markovian case is discussed as a special case in Subsection 3.9, including definitions (2.1') and (3.4') which we state here as (V.8) and (V.9), and an alternative characterization of the Gittins index in (3.15). While the authors consider non-negative utilities that are not necessarily bounded from above, it

(1994) and make use of the *retirement formulation* due to Whittle (1980). Here, the Gittins index is given by

$$M_i(t) = \inf\{m \in \mathbb{R} : L_i(c_i(t), m) = m\}, \quad (\text{V.8})$$

where we define the function  $L_i : \mathbb{R}^2 \rightarrow \mathbb{R}$  as

$$L_i(c_i(t), m) := \sup_{\tau \geq 0} \mathbb{E}_{t, c_i(t)} \left[ \int_0^\tau e^{-\alpha s} u_i(c_i(t+s)) ds + m e^{-\alpha \tau} \right]. \quad (\text{V.9})$$

The supremum is taken over all stopping times  $\tau$ . Appendix C.2 provides more details on this definition. The function  $L_i$  is the solution to an optimal stopping problem where we collect the discounted reward  $u(X)$  until a stopping time  $\tau$ . At  $\tau$ , we collect the retirement reward  $m$ . The Gittins index is now the smallest retirement reward  $m^*$  for which we are willing to stop immediately. In some sense, the Gittins index corresponds to the fair option price of continuously investing in one technology. The Gittins index strategy then *always invests in the technology with the highest option price*.

Since the  $M_i(t)$  are continuous variables, and technology costs constantly change upon investments, the Gittins indices of two technologies are essentially always distinct. This means that the Gittins index strategy essentially never invests in two technologies simultaneously (or even never if we choose a priority scheme). This is a crucial result. Similar to a static portfolio without risk aversion (Way et al., 2019), the dynamic portfolio also does not diversify its investments. However, unlike in the static case, the dynamic portfolio changes its investments over time if the expected utility is not realized. The opportunity to dynamically adjust investments in the future allows one to take on more short-term risks. As discussed in Section V.5, this result is notably different to what we observe in practice. Due to feedback delays, switching costs, risk aversion (non-linear utility functions) and heterogeneous preferences, investment portfolios are usually diversified, even if this is not theoretically optimal under the assumptions taken here.

The Gittins index strategy also has important computational implications. The fact is straightforward to adapt their results to the case of non-positive utilities.

that the indices  $M_i$  are independent of each other means that the computational effort of this strategy scales linearly in the number of technologies. Since the strategy only invests in one technology at a time, the  $M_i$ s only change for the technology into which one invests. This is a significant computational advantage, as our strategy only updates one  $M_i$  at a time.

The challenge with the Gittins index strategy is that the  $M_i$ s are generally difficult to compute. With the exception of some types of stochastic processes, such as stationary diffusion processes (Karatzas, 1984) or Levy-diffusion processes (Kaspi and Mandelbaum, 1995), there is no analytical solution for  $M_i$  (Powell, 2011). In the next two subsections, we show how  $M_i$  can be characterized analytically and then computed numerically.

### V.3.2 Gittins Indices for Innovation Bandits

Since the Gittins index  $M_i$  is calculated separately for each portfolio choice  $i = 1, \dots, d$ , we shorten our notation by dropping the subscript  $i$  for the remainder of this section.

To calculate the Gittins index, we can follow two steps. First, we solve the optimal stopping problem, Equation (V.9). Then, we use the solution  $L$  to solve for the index  $M$ , Equation (V.8). The main challenge lies in the optimal stopping problem (V.9). We follow the approach by Karatzas (1984) and translate it into a *free boundary problem*, where for all retirement rewards, we find the free boundary  $b_m(t)$ . This solves the optimal stopping problem in the sense that we stop the cost process iff  $c(t) > b_m(t)$ . Furthermore, the free boundary directly relates to the Gittins index based on the following theorem.

**Theorem 1.** *The Gittins index of an innovation bandit at time  $t$  and cost  $c$  is given by*

$$M(t; c) = \inf\{m \in \mathbb{R} : b_m(t) = c\}, \quad (\text{V.10})$$

where  $b_m(t)$  is the unique solution to the integral equation

$$\int_t^\infty e^{-\alpha s} \mathbb{E}_{t, b_m(t)} [(u(c(s)) - m\alpha) \cdot \mathbf{1}_{c(s) > b_m(s)}] ds = 0, \quad (\text{V.11})$$

for all  $t \in [0, \infty)$ .

*Proof.* The proof of this theorem is based on formulating the Gittins index calculation as a free-boundary problem. See [C.2.2](#).  $\square$

Based on this theorem, we need to find the retirement utility  $m^*$  for which  $b_m(t)$  falls on the current technology cost. Once we know  $b_m(t)$ , this greatly simplifies the search for the Gittins index through the map  $m \mapsto b_m(t)$ . Since Equation [\(V.11\)](#) can be computationally demanding to solve for individual values of  $m$ , it is helpful to apply the following translational invariance:

**Corollary 2.** *If the utility function  $u(t)$  fulfills  $a \cdot u(c) = u(c + f(a))$ , then for any  $m > 0$ ,*

$$b_m(t) = b_1(t) + f(m^{-1}). \quad (\text{V.12})$$

*Proof.* See Appendix [C.2.2](#).  $\square$

In our case, where  $u(c) = -e^c$ , this corollary implies that  $b_m(t) = b_1(t) - \log(m)$ . Now, instead of having to calculate the free boundary  $b_m(t)$  for a given value of  $c$  and solving for  $m$ , we can split this computation into two steps. First, we calculate  $b_1(t)$  for all values of  $t \geq 0$ . Then, for a given value of  $c$ , we use Corollary 2 to find  $M(t; c)$  for all values of  $t \geq 0$ . This effectively allows us to perform all resource-intensive calculations in pre-processing and find the Gittins index via a simple closed-form equation in each simulation.

### V.3.3 Numerical Solutions to the Gittins Index Strategy

To calculate the free boundary  $b_1(t)$  numerically, we follow a straightforward backward-integration method, similar to [Pedersen and Peskir \(2022\)](#), [Hou et al. \(2000\)](#), and [Huang et al. \(1996\)](#).

We first note that as  $t \rightarrow \infty$ ,

$$b_1(t) \rightarrow \log \left( \alpha \left( 1 + \frac{\sigma}{\sqrt{2\alpha}} \right) \right) =: b_1(\infty) \quad (\text{V.13})$$

based on Karatzas (1984), since the drift  $\mu(t)$  vanishes over time. This means that, in our numerical approximation, we can set  $b_1(T) \approx b_1(\infty)$  for sufficiently large  $T$ .

We then consider a time-interval  $[0, T]$  with large  $T > 0$ . We split this interval into  $n - 1$  disjoint sub-sections and consider times  $t_1 < t_2 < \dots < t_i < \dots < t_n$ . Since the free boundary depends exclusively on future values of  $b_1$ , we can apply backwards induction to find the step-function approximation  $\tilde{b}_1(t)$  of  $b_1(t)$ , where

$$\tilde{b}_1(t) = \begin{cases} \tilde{b}_1^i & \text{for } t \in [t_{i-1}, t_i) \\ b_1(\infty) & \text{for } t \geq T. \end{cases} \quad (\text{V.14})$$

For now, we consider  $\tilde{b}_1(t) \equiv b_1(t)$  for  $t \geq T$ , but we relax this condition later for practical purposes.

We apply Equation V.11 to find a recursive formulation of the free boundary. For ease of notation, we write  $b \equiv b_1$  and

$$J_b(t, c, s) \equiv \mathbb{E}_{t,c} [(u(c(s)) - \alpha) \mathbf{1}_{c(s) > b(s)}].$$

In our setting of linear utility and experience curves, this can be written explicitly as

$$J_b(t, c, s) = e^{c+(s-t)\sigma^2/2} \left( \frac{Z(s)}{Z(t)} \right)^{-\omega} \operatorname{erfc} \left( \frac{c - b + (s - t)\sigma^2 - \omega \log(\frac{Z(s)}{Z(t)})}{\sigma \sqrt{2(s - t)}} \right) - \alpha \operatorname{erfc} \left( \frac{c - b - \omega \log(\frac{Z(s)}{Z(t)})}{\sigma \sqrt{2(s - t)}} \right),$$

where we write  $b \equiv b(s)$  to shorten notation and  $\operatorname{erfc}$  corresponds to the complementary error function.<sup>9</sup> Provided  $(\tilde{b}_1^i)_{i>1}$  are known, the free boundary  $\tilde{b}_1^1$  is then given by

$$\sum_{i=1}^{n-1} \int_{t_i}^{t_{i+1}} e^{-\alpha s} J_{\tilde{b}_1^{i+1}}(t_0, \tilde{b}_1^1, s) ds + \int_T^\infty e^{-\alpha s} J_{b_1}(t_1, \tilde{b}_1^1, s) ds = 0. \quad (\text{V.15})$$

There are different options for our choice of  $t_i$  in this setting. Since the change in drift  $d\mu(t)/dt$  is largest for small  $t$ , we also expect the change in the free boundary  $db(t)/dt$  to be largest for small  $t$ . Since we are approximating  $b(t)$  with a step function, we use a

---

<sup>9</sup>The complementary error function on the real line is defined as  $\operatorname{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-s^2} ds$ .

smaller step size for small  $t$ . We choose  $t_i$  to be exponentially distributed to reflect this non-linearity. This improves both computation time and accuracy of our free boundary estimation compared to equidistant  $t_i$ . To obtain the free boundary for other values of  $t$ , we can interpolate  $\tilde{b}_1^i$  at  $t_i$ .

The resulting free boundary  $\tilde{b}_1(t_i)$  provides a first-order approximation of the restricted optimal stopping problem, where we are only allowed to make investment decisions at times  $t_i$ .  $\tilde{b}_1(t_i)$  converges to the true free boundary for  $n \rightarrow \infty$  and  $T \rightarrow \infty$  for both the continuous-time problem and the restricted problem. This convergence is unique and does not depend on the order of  $n, T \rightarrow \infty$  by the dominated convergence theorem.<sup>10</sup> For Wiener processes (without drift), the continuous free boundary converges to the free boundary of the restricted problem with  $\mathcal{O}(\sqrt{1/n})$  (Chernoff, 1965; Yao, 2006). While convergence speed and potential second-order corrections are non-trivial and out of the scope for this paper, the following examples show that the resulting Gittins index strategy still performs well. Moreover, Katehakis and Veinott (1987) show that an approximate index leads to an approximate optimal strategy for finite  $n$  and  $T$ . In other words, the numerical approximation of the free boundary still provides a strategy that is close to optimal.

Figure V.1 shows the resulting free boundary  $b_1$  for technology bandits of different learning parameters, initial cumulative deployment  $Z(0) = 1000$ , annual demand  $D = 1000$ , and  $\alpha = 0.05$ . The x-axis shows time  $t$ . The colors indicate different values for the learning parameter  $\omega$ . The line style shows different standard deviations  $\sigma$  of the noise. We can see that the boundary is convex and reaches  $b_1(\infty)$  asymptotically. The slope of the boundary depends on  $\omega$ , with fast-learning technologies showing a steeper boundary. In the  $\omega \rightarrow 0$  limit, the boundary is constant through time. We also see that the noise parameter  $\sigma$  has a large effect, with larger boundaries for higher noise levels.

The fact that the free boundary increases with volatility  $\sigma$  means that the investor is willing to bet on a volatile technology with large  $\sigma$  since they can divest if the associated

---

<sup>10</sup>The step function  $J_{\tilde{b}_1^{i+1}}(t_0, \tilde{b}_1^i, s)$  is bounded, such that we can apply the dominated convergence theorem to Equation (V.15). This means that we can exchange the limits of  $n \rightarrow \infty$  and  $T \rightarrow \infty$ .




Figure V.1: **Free boundary  $b_1(t)$  for different parameter values  $\omega$  and  $\sigma$ .** We compute the free boundary  $b_1(t)$  for different values of  $\omega$  and  $\sigma$  using the backward integration method. Since we assume  $Z(1) = 1000$  and linear growth, the time parameter is equivalent to the relative cumulative deployment. We set  $\alpha = 0.05$ . We see that the boundary is higher for larger values of  $\sigma$  and  $\omega$ . Due to the non-linear cost decline, the boundary is steeper for larger values of  $\omega$ . Towards  $t \rightarrow \infty$ , the learning exponent  $\omega$  becomes irrelevant.

risk does not pay off. In this sense, their behavior is volatility seeking. This is the dominant behavior in the absence of learning (i.e. large  $Z(t)$  or small  $\omega$ ) and the investor favor technologies with large  $\sigma$ . In the presence of learning (i.e. small  $Z(t)$  or large  $\omega$ ), there is a trade-off between learning  $\omega$  and volatility  $\sigma$ —the investor may choose a technology with higher  $\omega$  and lower  $\sigma$ . Note that, since investment decisions also depend on current costs, lower free boundary parameters relative to others do not necessarily mean that the investor will only invest in that technology.

### V.3.4 Gittins Index vs. Expected Net Present Cost

An alternative method commonly applied in the context of allocating resources between technologies is based on expected net present cost (ENPC) (Mazzola and McCardle, 1996). In the static setting, the optimal strategy is to calculate the ENPC for each technology

and invest all resources into the technology with the smallest ENPC (Way et al., 2019). In the dynamic setting, this also constitutes an index strategy, albeit not a theoretically optimal one. For initial cost  $C(t) > 0$  and deployment  $Z(t)$  at time  $t > 0$ , the ENPC  $N(t)$  is given by

$$N(t) = \mathbb{E}_{t, C(t)} \left[ \int_0^\infty e^{-\alpha s} C(t+s) ds \right] = C(t) \int_0^\infty e^{-\gamma s} \left( \frac{Z(t)}{Z(t) + Ds} \right)^\omega ds \quad (\text{V.16})$$

$$= C(t) \left( \frac{Z(t)}{D} \right)^\omega \gamma^{\omega-1} e^{Z(t)\gamma/D} \Gamma \left( 1 - \omega, \frac{Z(t)\gamma}{D} \right), \quad (\text{V.17})$$

where  $\gamma = \alpha - \sigma^2/2$  (for a shorter notation),  $\Gamma(\cdot, \cdot)$  is the incomplete Gamma function, and  $\sigma^2 < 2\alpha$  by assumption. This equation follows from Fubini's theorem.

---

**Algorithm 2:** Expected Net Present Cost (ENPC) index strategy

---

For each technology  $1 \leq i \leq d$  at time  $t_0 \geq 0$ , calculate the ENPC  $N_i(t) \in \mathbb{R}$ .

**while**  $t > t_0$  **do**

Find technology  $j$ , with  $N_j(t) = \min_i \{N_i(t)\}$ .

**while**  $N_j(t) = \min_i \{N_i(t)\}$  **do**

Invest in technology  $j$  i.e.  $\dot{\zeta}_j(t) > 0$

Update the ENPC  $N_j(t)$

---

To explore the difference between the Gittins index strategy and the ENPC index strategy<sup>11</sup>, we compute both for two technologies (A and B) using different discount rates and relative learning parameter values. The results are presented in Figure V.2. We set  $\sigma_A = 0.15$  and  $\omega_A = 0.15$  and consider different relative values for  $\sigma_B$  and  $\omega_B$ . The initial cost is equal to 1 in both cases. The orange area highlights parameter values where the Gittins index strategy would initially invest into technology A (and B otherwise). The striped area represents values where the ENPC index strategy would initially invest into technology A (and B otherwise).

The difference in strategies is striking. If we set  $\sigma_B = \sigma_A$  and vary  $\omega_B$ , both strategies

---

<sup>11</sup>In our dynamic allocation problem, we consider negative costs,  $-C(t)$ , as utility so that a strategy aims at maximizing utility. For the ENPC index strategy, we use the equivalent perspective of aiming to minimize costs.

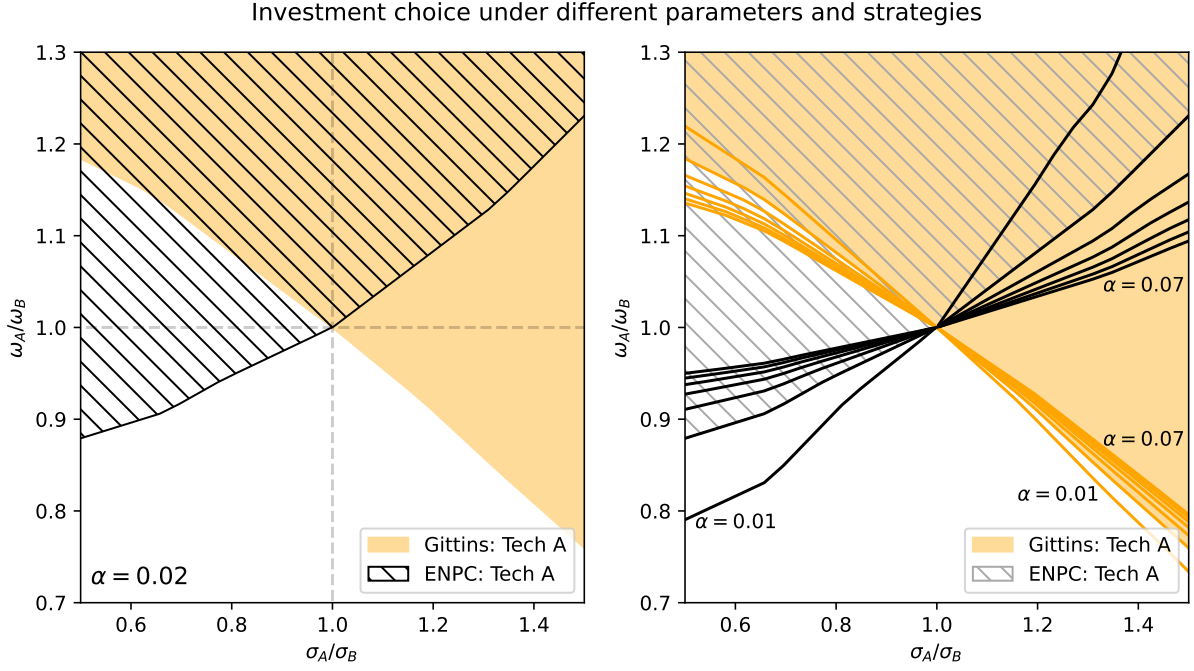


Figure V.2: **Comparison of the initial investment choice under the Gittins and ENPC strategy for different parameter values.** We consider two technologies  $i \in \{A, B\}$  with Wright's law parameters  $\omega_i, \sigma_i$  and initial cost  $c_i = 1$ . For different relative values of  $\sigma$  (x-axis) and  $\omega$  (y-axis), we compute the initial investment choice based on the Gittins index and ENPC of each technology. In the orange area, the Gittins strategy invests into technology A, in the white area into technology B. In the striped area, the ENPC strategy invests into technology A and B otherwise. **(Left)** For the discount factor  $\alpha = 0.02$ , we see that both strategies differ substantially. While both prefer higher learning rates  $\omega_A$ , the Gittins strategy favors the technology with higher noise  $\sigma$ , while the ENPC strategy favors lower noise levels. **(Right)** This preference holds for different discount factors. However, the non-linear trade-off between  $\omega$  and  $\sigma$  is sensitive to  $\alpha$ , particularly in the ENPC strategy. For low discount factors, the noise parameter  $\sigma$  becomes more important to the investment decision.

invest in technology B for  $\omega_B > \omega_A$ , and technology A otherwise. This is not surprising since more learning is associated with larger cost reductions. However, if  $\omega_B = \omega_A$  and  $\sigma_B > \sigma_A$ , the Gittins strategy will invest in B while the ENPC strategy invests in A (and vice-versa for  $\sigma_B < \sigma_A$ ). Investing in technologies with larger volatility is optimal in the dynamic setting even though the ENPC is higher. This is because higher values of  $\sigma$  lead to higher ENPC but also to an increased chance of cost reduction, making it more attractive in a dynamic setting. The fact that the decision boundary is not linear in either strategy indicates that there is a non-linear trade-off between gaining experience and volatility in both cases.

In addition to the experience curve parameters, the investment choice between technologies A and B also depends on the discount rate  $\alpha$ . We investigate this relationship in

the right panel of Figure V.2. For low values of  $\alpha$ , the decision boundaries for the ENPC and Gittins strategy become more vertical. A low discount factor implies that technology costs in the distant future contribute similarly to the net present cost (NPC) as technology costs in the near future. For  $\alpha \rightarrow 0$ , both decision boundaries become independent of  $\omega$ , since only events in the distant future without learning effects are relevant for the NPC. Conversely, for  $\alpha \rightarrow \infty$ , the decision boundaries become horizontal (i.e., independent of  $\sigma$ ). In this limit, only immediate events impact the discounted cost where the effect of the noise vanishes.

Finally, initial costs,  $c_0$ , also influence the investment decisions. The initial cost is a multiplicative factor for both ENPC and Gittins strategy, such that both strategies scale equally. We will explore the role of initial costs more closely in our empirical example of light passenger vehicle investments.

### V.3.5 Gittins Index for Mean-Reverting Bandits

While many technologies can be described according to Wright’s law, mean-reverting processes are better suited for some. For example, fossil fuel technology costs have been remarkably stable over long periods in recent history with just some short-term noise around the mean (Way et al., 2022; McNerney et al., 2011a).

In our continuous time setting, we can describe these technologies using an Ornstein-Uhlenbeck process (Way et al., 2022), where the logarithmic cost  $\log C(t) =: c(t)$  solves

$$dc(t) = \theta(\kappa - c(t))dt + \sigma dB(t) \tag{V.18}$$

in Itô sense. Again,  $B$  is a standard Brownian motion started from 0,  $\sigma > 0$  is the noise parameter,  $\theta > 0$  is the autocorrelation parameter, and  $\kappa \in \mathbb{R}$  is a constant defining the long-term average of  $c(t)$ .

Theorem 1 cannot be used to calculate the Gittins index for an Ornstein-Uhlenbeck process since the drift is stochastic and not monotone. However, since the Ornstein-

Uhlenbeck process is time-homogeneous, the Gittins index can be computed based on [Karatzas \(1984, Theorem 3.1\)](#). We refer the reader to [Appendix C.2.1](#) for a detailed computation leading to the following corollary.

**Corollary 3.** *The Gittins index  $M$  of a mean-reverting bandit at time  $t$  and log cost  $c$  can be calculated as*

$$M(c) = -\frac{\sigma\sqrt{\theta}}{2\alpha} \frac{H_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - c) \right)}{H_{-\alpha/\theta-1} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - c) \right)} p'(c) + p(c), \quad (\text{V.19})$$

where  $H_\nu(\cdot)$  is the probabilistic Hermite polynomial of order  $\nu$  and

$$p(c) = -\mathbb{E}_c \left[ \int_0^\infty e^{-\alpha t} e^{c(t)} dt \right] = -\int_0^\infty e^{-\alpha t} e^{ce^{-\theta t} + \kappa(1-e^{-\theta t})} \cdot \mathbb{E}_0 \left[ e^{\sigma \int_0^t e^{-\theta(t-s)} dB(s)} \right] dt. \quad (\text{V.20})$$

This corollary allows us to calculate the Gittins index  $M$  for mean-reverting bandits given the Hermite polynomials  $H$  and integral function  $p$ . While neither have a simple close-form representation, there are efficient numerical schemes to estimate  $H$ . For  $p$ , we can simulate the stochastic process  $c(t)$  to obtain a numerical approximation. Fitting this approximation with a high-degree polynomial yields an efficient numerical estimates to both  $p(c)$  and  $p'(c)$ .

## V.4 Application to Light Vehicle Investments

We now turn to an empirical example in the transportation sector. Transitioning light passenger vehicles from fossil-fuel-based internal combustion engines (ICEs) to battery-electric vehicles (BEVs) is imperative to the energy transition ([IPCC, 2022](#)), but BEVs still face higher upfront capital costs than ICEs in most parts of the world ([IEA, 2024](#)). We model different dynamic investment strategies to illustrate the Gittins strategy and better understand how managers, investors, and policymakers can navigate this transition. We focus on the U.S. light passenger vehicle market due to data availability.

### V.4.1 Light Vehicle Experience Curves

For the remainder of this section, we consider BEV and ICE prices. Although BEV prices are currently higher than ICE prices, they are subject to an experience curve. In particular, battery packs, which make up a large share of the total EV price, have been subject to significant declines (Way et al., 2022; Weiss et al., 2019; McKerracher, 2024; Kittner et al., 2020). On the other hand, ICE prices have not shown any significant trend since the 1970s and even increased between 1970 and 2000 (DOE, 2021).

In reality, consumer decisions will depend not just on the vehicle price, but also fixed and variable operating costs such as expenditure for petrol and electricity. We exclude these costs for the presented example since they are impacted by exogenous market conditions rather than endogenous innovation dynamics. In the future, alternative index strategies, such as the Whittle index, can be used to capture these dynamics more holistically, see section V.5.3.

Figure V.3 shows the historical prices of light vehicles in the U.S., with black dots representing average ICE prices (DOE, 2021). Prices have been approximately constant since the 1990s. Green dots capture (modelled) representative BEV prices in the U.S. Since the BEV market has historically focused on premium market segments, we model BEV prices based on historic battery costs comparable to recent mid-segment ICE vehicles. We see that BEV prices declined by an order of magnitude since the 1990s. At the same time, the stock of BEVs and plug-in hybrid vehicles (PHEVs) has increased by over 3 orders of magnitude since 2005 (represented by the green solid line). Future BEV prices may decline further if BEV sales continue growing due to experience, while ICE prices are expected to remain relatively constant.

To compare future BEV and ICE prices  $\mathcal{P}_i$ , we distinguish between base prices  $\mathcal{B}$  and BEV/ICE-specific powertrain prices  $\mathcal{M}_i$ , such that

$$\mathcal{P}_i = \mathcal{B} + \mathcal{M}_i, \tag{V.21}$$




Figure V.3: **Historical stock and prices of BEVs and ICEs in the US.** Since the 1990s, BEVs (green) have declined in cost by an order of magnitude and grown in stock by almost 3 orders of magnitude. Over the same period, ICE prices have remained approximately constant. ICE prices are based on [DOE \(2021\)](#). BEV prices are modelled based on historical Li-Ion battery costs from ([Way et al., 2022](#)), assumed battery capacity of 73kWh, and price intercept from [DOE \(2022\)](#). BEV + PHEV stock is taken as the combined battery-electric and hybrid-electric stock from [DOE \(2024a,b\)](#).

where  $i \in \{BEV, ICE\}$ . The base cost  $\mathcal{B} \approx 26,000$  USD is fixed over time. It includes components such as the chassis and interior that are highly comparable between powertrain technologies ([DOE, 2022](#); [König et al., 2021](#)).<sup>12</sup> We focus on modeling powertrain prices since shared components do not impact the decision between the two technology types. For BEVs, the powertrain price  $\mathcal{M}_{BEV} \approx 10,000$  USD is based on a 73 kWh Li-Ion battery at  $\sim 140$  USD/kWh in 2023 since it is the largest cost component ([König et al., 2021](#)). For ICEs, the implied powertrain price  $\mathcal{M}_{ICE} \approx 2,500$  is based on Equation V.21 and current price  $\mathcal{P}_{ICE} \approx 28,500$  USD.

For the  $\mathcal{M}_{BEV}$ , we use historic battery costs to estimate the experience curve and predict future prices. We apply the parameters shown in Table V.1, which are estimated by [Way et al. \(2022\)](#) using the method developed in [Lafond et al. \(2018\)](#). Each doubling

<sup>12</sup>BEV/ICE powertrains can also contain comparable components, such as differentials. We account for these components in the base price  $\mathcal{B}$ .

Parameter	BEVs	Parameter	ICEs
$\omega$	0.421	$\theta$	0.9
$\sigma$	0.103	$\kappa$	7.7
		$\sigma$	0.1

Table V.1: **Experience curve and mean-reversion parameters for BEV and ICE prices.** We calibrate the experience curves using the method of Lafond et al. (2018). We calibrate the mean-reversion parameters using a maximum likelihood estimator.

in cumulative EV sales (including hybrid EVs) leads to a  $1 - 2^\omega \approx 25\%$  drop in median battery costs.

For the  $\mathcal{M}_{ICE}$ , we assume a mean-reverting cost model in line with Way et al. (2022) and McNerney et al. (2011a), and consistent with how costs for ICE vehicles have been approximately constant over the last 30 years (after adjusting for inflation). We derive the parameters of the Ornstein-Uhlenbeck process using a maximum likelihood estimator (see Table V.1).

In addition to the experience curve parameters, our strategies also depend on annual demand. Light vehicle sales in the U.S. grew between 1970 and 2000 but have since stagnated (DOE, February 6, 2023). We, therefore, assume that light vehicle demand will remain constant and apply an annual demand of 15.6 million vehicles based on the average number of units sold since 2000.

We can now calculate the Gittins index according to the free-boundary estimation developed in Section V.3 for BEVs, and for ICEs, we rely on Corollary 3.

## V.4.2 Light Vehicle Investment Strategy

Using these experience curves as well as current technology prices and cumulative deployment, we simulate different investment strategies for a BEV/ICE product portfolio under the Gittins index strategy (Algorithm 1) versus the ENPC index strategy (Algorithm 2). Figure V.4 shows the results in terms of the NPC of each strategy (left panel) and the underlying portfolio choices (right panel).

We find striking differences. Following the ENPC strategy would lead to investing all resources in ICEs, which is driven by the high prices of BEVs, their associated uncertainty,

and the relatively low volatility of ICE prices under the mean-reverting model. Under this strategy, BEV prices remain unchanged and ICE prices fluctuate around their long-term mean, although not strongly enough to warrant an investment in BEVs. As a result, there is almost no variation in the strategy's NPC of approximately 26,000 USD.

The Gittins strategy, on the other hand, initially invests all resources in BEVs due to the possibility of future price reductions. However, if these price reductions do not materialize, the Gittins strategy switches its investments to ICEs. Around 40% of simulation runs indicate investment switching after 140 years. As a result, the NPC of the Gittins strategy displays a bimodal distribution with significant variation between 0.5tn USD and 4tn USD. The mean NPC of the Gittins strategy at 2.4tn USD is around 7% below the mean of the ENPC strategy.

It is important to note that this result is sensitive to the applied discount rate. Here, we have chosen a discount rate  $\alpha = 1\%$  that a social planner may use. An automotive manager will likely use a discount rate that is significantly higher (double-digit). In this case, neither Gittins nor ENPC strategy initially invest into BEVs since they weight the lower costs of ICEs higher than the future BEV cost reductions. This is similar to [Way et al. \(2019\)](#), where small changes in parametrization can lead to substantially different results.

Neither strategy shows large volatility in the investment behavior; The ENPC strategy consistently invest in ICEs, the Gittins strategy may or may not switch once between technologies. This suggests that there are different tipping points associated with each strategy. [Figure V.5](#) shows the price BEVs would have to reach for the Gittins and ENPC algorithms to place their initial investment into them, depending on cumulative deployment. One invests in BEVs at a higher powertrain price (and thus sooner) under the Gittins strategy relative to the ENPC strategy due to how uncertainty is treated. The decision boundaries are approximately parallel, although they start to diverge when cumulative deployment is high.

[Figure V.5](#) also shows historical powertrain prices for BEVs and ICEs. BEV powertrain

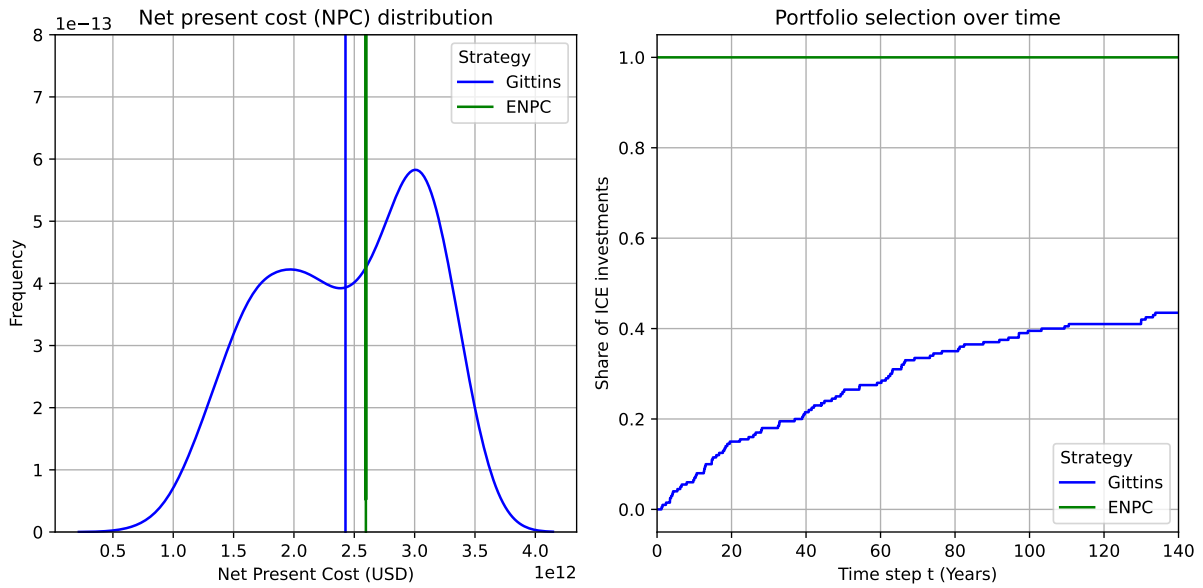


Figure V.4: **Net Present Cost and portfolio selection of the Gittins and ENPC index strategies at  $\alpha = 1\%$  discounting.** We simulate the Gittins and ENPC strategies for ICE/BEV portfolios. **(Left)** NPC distribution under the Gittins and ENPC strategies. The ENPC strategy shows almost no variation in its NPC since it invests exclusively in ICEs. The Gittins strategy shows significant bimodal variation in NPC. The mean NPC, indicated with a vertical line, is 7% smaller than under the ENPC strategy. **(Right)** The ENPC strategy invests exclusively in ICE vehicles. The Gittins strategy invests initially in BEVs. In later time steps, the strategy shifts increasingly towards ICEs. After 140 steps, around 40% of simulation runs invest in ICEs.

prices decline with cumulative deployment, along the power-law relationship exhibited by experience curves. In recent years, prices declined faster than expected by the experience curve due to their uncertainty. As a result, if we fix the experience curve to the BEV powertrain price in 2023, prices in earlier years appear above the long-term trend. Before 2021, BEV powertrain prices are above both decision boundaries such that neither strategy would have (ever) invested in BEVs due to lock-in effects. BEV prices fall low enough in 2021-2023 for the Gittins strategy to cross its tipping point and invest in BEVs, though, whereas those following an ENPC strategy never would invest in BEVs.

It is striking that the decision boundaries and experience curve run approximately parallel in the current cumulative deployment regime. As BEV powertrain prices decrease along their experience curves, they will most likely not get closer to the decision boundary of either Gittins or ENPC strategy.

The simplicity of this result shows that reality is usually more complex than cost-

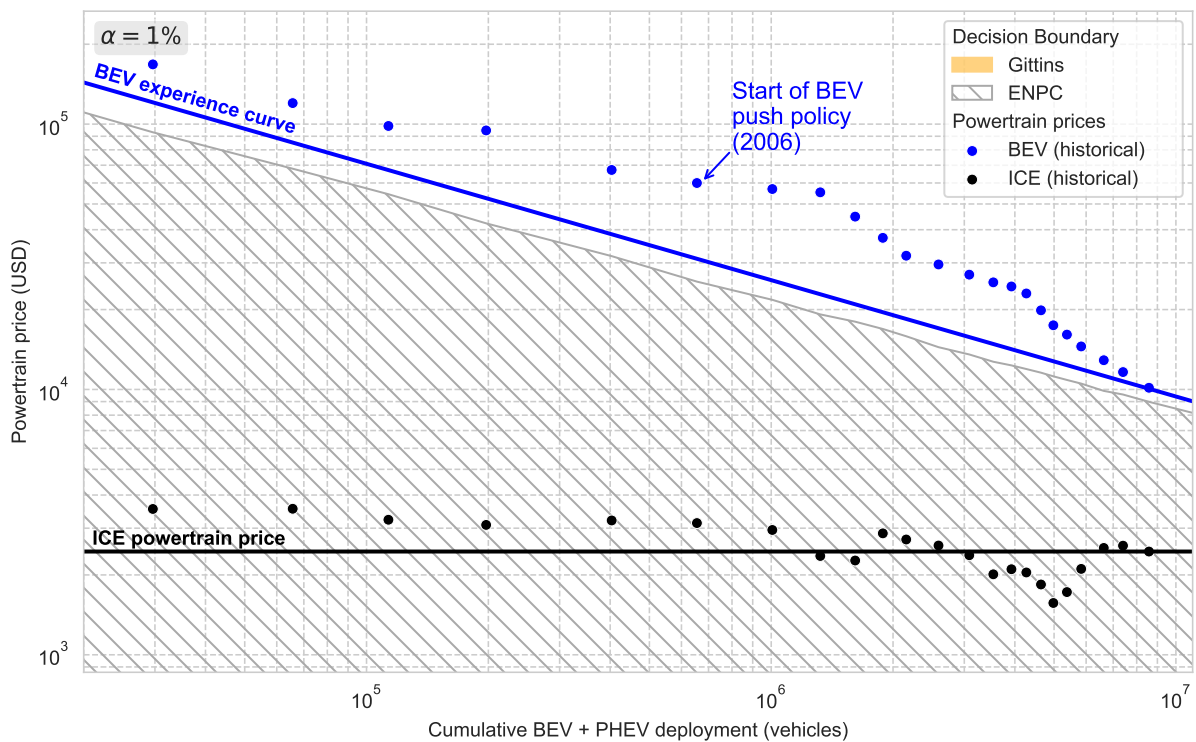


Figure V.5: **Gittins and ENPC decision boundary against cumulative EV deployment.** (Dots) BEV and ICE powertrain prices against the cumulative deployment of electric vehicles in the US. Both prices and deployment are shown on a logarithmic scale. ICE powertrain prices can be modeled with a mean-reverting AR(1) process around the long-term average, shown in black. BEV powertrain prices are decreasing along the experience curve, shown in blue. **(Orange area)** Investment decision boundaries under the Gittins strategy. Within the orange area, the Gittins strategy invests in BEVs, outside it invests ICEs. **(Grey area)** Investment decision boundaries under the ENPC strategy. Within the grey-striped area, the ENPC strategy invests in BEVs, outside it invests ICEs. Up until 2021, both algorithms invested exclusively into ICEs. However, with decreasing BEV powertrain prices, the Gittins algorithm now initially invests in BEVs. The BEV investment prior to 2021 was in part driven by the adoption of various technology-push policies after the Energy Policy Act of 2005 (DOE, n.d.).

optimisation. Next to maximising NPC, decision-makers must further account for policy and market signals, as well as potential externalities, such as those posed by climate change. In some cases, for example if BEVs were to follow their experience curve exactly, innovative technologies may require continued policy support despite strong price declines.

### V.4.3 Dynamic Policy Choices

Dynamic policy interventions can play a crucial role in accelerating technological change by both reducing costs and creating critical tipping points. Policies may be broadly classified into two types: those that reduce costs through a technology-push mechanism

(such as R&D grants, tax incentives, or regulatory support) and those that boost demand (demand-pull policies like consumer adoption subsidies or mandated purchase targets). Demand-pull policies not only stimulate market growth but also affect costs via experience curves.

In the U.S., the policy landscape for electric vehicles (EVs) has evolved through a combination of technology-push and demand-pull measures. The Energy Policy Act of 2005 marked one of the first major explicit demand-pull interventions by providing federal tax breaks for EV purchases. More recently, policies such as those embedded in the Inflation Reduction Act (IRA) have provided additional consumer subsidies. Since 2006, EV deployment has grown by about 20% annually on average despite BEV costs initially remaining above the ENPC and Gittins index decision boundaries.

One implicit goal of demand-pull policies is to move technologies along the experience curve until crossing a decision boundary. However, uncertainties in learning (and deployment) make the length of time that support is required unclear. To explore the dynamic implications of demand-pull policies, we simulate a scenario that mimics the policy environment beginning in 2006. We assume that a demand-pull policy boosts BEV deployment by 20% annually (similar to what has transpired in reality) and BEV prices evolve endogenously according to Wright’s law. Once costs fall below a decision boundary (either as determined by the Gittins index or the ENPC criterion), this signals a tipping point where policy support can be phased out.

Our results indicate that, for the Gittins index decision boundary, the expected time to reach the tipping point would have only been around 2035. However, the uncertainty of this estimate is large—with a standard deviation of more than 13 years. Because of innovation uncertainty, BEVs actually crossed the tipping point much earlier (2021) and are now close to cost parity with ICEs ([Schmidt et al., 2017](#)). The probability of never hitting the decision boundary is very low, with less than a 4% probability of needing to continue support beyond 2060. Conversely, the ENPC-based approach yields a much later tipping point, with an expected crossing around 2049 (and a similar degree of uncertainty)

along with a considerably higher probability (over 37%) that policy support would need to continue beyond 2060. These results have several important implications for future policy design, which we discuss in Section [V.5](#).

## V.5 Discussion

Our analysis demonstrates that MAB-based optimization strategies can be applied to innovation and technology portfolios, with advantages over conventional approaches based on a static framework. We calculate one such strategy based on the Gittins index explicitly and demonstrate its use on an example of electric vehicle investments. This also highlights the strong assumptions and limitations of the Gittins approach. In this section, we provide further discussion of investment strategies in this setup, implications for public policy, and opportunities for extensions of the MAB framework in the context of innovation.

### V.5.1 Dynamically Optimal Innovation Portfolios

By reformulating the innovation portfolio problem within a MAB framework, we show that optimal investment and planning decisions can be volatility-seeking and non-diversified under certain conditions, particularly low switching costs and fast feedback mechanisms.

A key result of the Gittins strategy is that the optimal dynamic portfolio concentrates investments into a single assets, thus averting diversification. The rapid feedback mechanism allows decision-makers to maximise the exploration of a single technology, with limited downside risk. Low switching costs allow the decision-maker to redeem poor investment choices by switching to alternatives with more positive outlook. As a result, it is efficient to focus on individual technologies and switch if they do not work out. This means that the Gittins strategy is best applicable to technologies with short development cycles and fast feedback, and less so if investment feedback is slow. For example, solar photovoltaics projects can be built within months, while nuclear projects can take over a decade. (This analogy is also imperfect due to supply chain constraints.) Application of a

fast-feedback strategy requires careful investigation of the underlying portfolio options. The relevant timescales further depend on the discount rate applied to the optimization problem.

The fact that decision-makers can redeem poor investment choices is especially relevant when costs are highly volatile. The resulting optimal portfolio choice is one with higher volatility even when volatility decreases the long-term expected return. This outcome is closely related to the assumed absence of risk-aversion in our utility function. Risk-aversion implies non-linear utility functions, which, in turn, usually imply diversified portfolios (Way et al., 2019). Moreover, a set of investors may have different risk aversion. This means that the tipping points, or decision boundaries, for different investors differ, and the overarching multi-investor portfolio is usually diversified and much less volatility-seeking than in the Gittins strategy.

## V.5.2 Public Policy Implications

Despite these limitations, our findings have implications for public policy and planning. First, our findings on non-diversified portfolios indicate that investments concentrated in a few promising technologies—especially those with relatively low switching costs—can be effective to rapidly identify winners and benefit from early learning effects. This stands in contrast to “neutral” policies that do not discriminate between portfolios (Fabra and Montero, 2023; de Mello Santana, 2016). Moreover, our conclusion that investing in portfolios with higher volatility can be optimal under certain conditions is consistent with Rodrik (2014)’s point that well-designed industrial policies inevitably come with mistakes and that too few mistakes may even signal *under*performance. In this sense, a dynamic, volatility-seeking strategy that tolerates short-term failures may ultimately enable faster cost reductions and technological transitions if our conditions on rapid feedback and low switching costs are met.

Our work also provides practical insight for designing adaptive policy and planning frameworks. In the context of early-stage deployment, policymakers may wish to consider

implementing flexible interventions that can be frequently reassessed as market conditions evolve. This adaptive approach can allow policymakers to provide support until a technology reaches its tipping point, thereby avoiding unnecessary public expenditures thereafter. The Gittins index strategy can help to prioritize technology options and dynamic simulations can be used to investigate the duration and cost-efficiency of associated policy support. This adaptive strategy contrasts to long-term planning scenarios. For example, the UK's Future Energy Scenarios (FES) lay out a narrow planning framework for the energy transition until 2050 that may be too restrictive to allow for adaptive planning of novel technologies (NESO, 2025).<sup>13</sup>

That said, implementing volatility-seeking dynamic strategies is likely challenging in practice, as expecting failure may be met with political resistance and there can be substantial benefits to planning certainty. Many technology portfolios also have significant switching costs, such as those associated with supply chain constraints. Nevertheless, having an optimal dynamic strategy can be helpful to better understand the full solution space of public funding strategies and the costs associated with inflexibility.

### V.5.3 Opportunities for Future Work

By demonstrating the application of MAB-based strategies in an empirical setting of dynamic, stochastic, and non-linear portfolio optimization, our work opens numerous avenues for future research. Extensions to the presented index-based solution could include incorporating parameter uncertainty, exploring portfolio constraints, and adapting the utility function to account for non-exponential discounting or risk aversion. These enhancements could further align the model with various real-world conditions, such as switching costs and feedback delays. There are also a number of non-parametric MAB strategies that do not rely on the Gittins index. These strategies are worth exploring in detail since they may require weaker underlying assumptions. Lastly, there are multiple open questions to identify when the assumptions taken in our approach and its alternatives

---

<sup>13</sup>For example, the FES lay out an very narrow offshore wind capacity range of 96.4-104.4 GW in 2050, compared to 15.5 GW installed capacity today.

are appropriate or not.

For the existing index-based solution space, there are several generalizations to consider with respect to the stochastic process of the technology choices. This includes generalizations towards stochastic processes with parameter uncertainty, such as the Wright's law exponent (e.g., [Yao \(2006\)](#) and [Cowan et al. \(2018\)](#)), the correlation between technologies (e.g., [Gupta et al. \(2021\)](#)), and restless technologies that may undergo evolution without investment (e.g., [Whittle \(1988\)](#); [Funk and Magee \(2015\)](#)). Another interesting extension could be bandit processes with delays, where investment choices may be required before discovering the result of previous investments to reduce our reliance on immediate feedback.

So far, portfolio utility has been limited to exponentially discounted, monotone, and linear functions. This can be extended to include switching costs ([Jun, 2004](#)) and risk aversion ([Huo and Fu, 2017](#); [Sani et al., 2013](#)). One may be particularly interested in considering [Malekipirbazari and Çavuş \(2024\)](#) for the latter, as they demonstrate that Whittle's retirement formulation can be extended to include dynamic risk measures in near-optimal index strategies. To some extent, this is similar to the previous generalization towards correlated portfolio choices since non-linear utility functions depend on investment choices at different points in time across (potentially) different choices. Non-exponential discounting also could be an important extension, with hyperbolic discounting being particularly useful when demand grows exponentially. It is also a better representation of the time preference observed in behavioral economics ([Fedus et al., 2019](#)).

Lastly, we are interested in generalized index-strategies under a set of constraints, particularly with respect to the previous extensions. This may include generalized strategies with constraints on the number of technologies in which one can invest in a given period. Another direction are portfolios of technologies that are not competing, but instead, complement each other towards achieving a common goal, or portfolios with multiple applications of the same technologies (e.g., [Sinsel et al. \(2020\)](#), [Chen et al. \(2013\)](#), and [Beuse et al. \(2020\)](#)). This could be particularly relevant for system-level optimization models, where a landscape of solutions is required to provide different services.

As we have seen here, index-based MAB solutions frequently require strong assumptions to make them computationally tractable. Alternatively, there are a number of non-parametric solutions to the MAB problem, such as Thompson sampling, upper-confidence bounds, and greedy strategies (Agrawal and Goyal, 2012; Gittins et al., 2011). These strategies are often simpler to implement than the Gittins index, show good heuristic performance under fewer constraints, and are increasingly well understood analytically (Scott, 2010). Examples from economics and operations research includes Shen et al. (2015) who use an upper confidence bound model for financial portfolio optimization, Kasy and Teytelboym (2023) applying a Thompson sampling variant to optimize employment in refugee resettlement, or Cheung et al. (2022) who apply a two-stage sliding window-upper confidence bound algorithm to pricing automotive loans. The latter is particularly relevant for innovation strategies since it includes both parameter uncertainty about dynamic switching costs (or budgets) and nonstationary drift.

Exploring these alternatives will likely require a re-formulation of the objective function. Most non-parametric strategies are set up to minimise the regret of their strategy with respect to a benchmark instead of utility-maximisation discussed so far (Lai and Robbins, 1985). This benchmark could either be a naive or static model serving as a lower bound, or an unobtainable upper-bound sometimes referred to as an oracle. The presented Gittins strategy is a natural choice of oracle in the present case of endogenous uncertainty since it is theoretically optimal. Additionally, applying non-parametric strategies in innovation will require extensive empirical or heuristic testing, as well as analytical performance bounds.

As a final point, investigating the limits of the different assumptions taken in any MAB-based solution is an important research direction in itself. As mentioned earlier, a quantitative understanding of where our approach is reasonable, and where not, will help identify the appropriate model for a given task. For example, understanding how much data on a technology is required to assume that the learning rate is known is an important research direction. The same goes for other model assumptions, such as the choice of objective function or constraints put on the solution space of dynamic portfolio strategies.

## V.6 Conclusion

In this paper, we develop a novel dynamic portfolio optimization framework that incorporates endogenous innovation using a Multi-Armed Bandit approach. We derive an integral formulation for the Gittins index under monotone, non-linear, time-inhomogenous drift and show how it can be solved numerically. Our numerical implementation—based on a backward-integration scheme—demonstrates both asymptotic optimality and computational efficiency.

We apply our strategy to different portfolios of technologies evolving under Wright’s law. Using an empirically validated stochastic learning curve model, we show that our approach outperforms commonly used strategies based on the ENPC of each technology, as dynamic strategies can adapt investment allocations in response to declining costs. The simulation results show that dynamic policies not only reduce net present costs but also help identify critical tipping points. These findings suggest that adaptive, volatility-seeking strategies may facilitate more rapid market transitions, providing important implications for both policymakers and private investors. At the same time, our approach is contingent on a number of strong assumptions, particularly low switching costs and fast investment feedback. We outline some modern MAB-based approaches to overcome these limitations and make our method more realistic in the future.

In summary, we provide a robust theoretical framework as well as a practical decision-making tool for managing innovation investments under uncertainty by extending existing MAB techniques to capture the endogenous and stochastic nature of technological innovation. Our work can help pave the way for more dynamic and responsive public and private investment strategies as well as open pathways for future applications in technology policy and strategic innovation management.

## References

- Agrawal, Shipra and Navin Goyal**, “Analysis of Thompson Sampling for the Multi-armed Bandit Problem,” *JMLR: Workshop and Conference Proceedings*, 2012, *23*, 39.1–39.26.
- Ballesterero, Enrique, Mila Bravo, Blanca Pérez-Gladish, Mar Arenas-Parra, and David Plà-Santamaria**, “Socially Responsible Investment: A Multicriteria Approach to Portfolio Selection Combining Ethical and Financial Objectives,” *European Journal of Operational Research*, January 2012, *216* (2), 487–494.
- Beuse, Martin, Bjarne Steffen, and Tobias S. Schmidt**, “Projecting the Competition between Energy-Storage Technologies in the Electricity Sector,” *Joule*, October 2020, *4* (10), 2162–2184.
- BNEF**, “BNEF Electric Vehicle Outlook 2023,” Technical Report 2023.
- Bolton, Patrick and Christopher Harris**, “Strategic Experimentation,” *Econometrica*, 1999, *67* (2), 349–374.
- Bouneffouf, Djallel, Irina Rish, and Charu Aggarwal**, “Survey on Applications of Multi-Armed and Contextual Bandits,” in “2020 IEEE Congress on Evolutionary Computation (CEC)” July 2020, pp. 1–8.
- Bukenberger, Jesse and Bryan Palmintier**, “Stochastic Generation Capacity Expansion Planning with Approximate Dynamic Programming,” *2018 IEEE/PES Transmission and Distribution Conference and Exposition (T&D)*, April 2018, pp. 1–5.
- Chao, Raul O. and Stelios Kavadias**, “A Theoretical Framework for Managing the NPD Portfolio: When and How to Use Strategic Buckets,” *Management Science*, May 2008, *54* (5), 907–921.

- Chen, Wei, Yajun Wang, and Yang Yuan**, “Combinatorial Multi-Armed Bandit: General Framework, Results and Applications,” *Proceedings of the 30th International Conference on Machine Learning*, 2013, 28.
- Chernoff, Herman**, “Sequential Tests for the Mean of a Normal Distribution IV (Discrete Case),” *The Annals of Mathematical Statistics*, February 1965, 36 (1), 55–68.
- Cheung, Wang Chi, David Simchi-Levi, and Ruihao Zhu**, “Hedging the Drift: Learning to Optimize Under Nonstationarity,” *Management Science*, March 2022, 68 (3), 1696–1713.
- Cowan, Robin**, “Tortoises and Hares: Choice Among Technologies of Unknown Merit,” *The Economic Journal*, 1991, 101 (407), 801–814.
- Cowan, Wesley, Junya Honda, and Michael N Katehakis**, “Normal Bandits of Unknown Means and Variances,” *Journal of Machine Learning Research*, 2018, 18, 1–28.
- de Angelis, Tiziano and G. Peskir**, “Global  $C^1$  Regularity of the Value Function in Optimal Stopping Problems,” *The Annals of Applied Probability*, June 2020, 30 (3).
- de Mello Santana, Paulo Henrique**, “Cost-Effectiveness as Energy Policy Mechanisms: The Paradox of Technology-Neutral and Technology-Specific Policies in the Short and Long Term,” *Renewable and Sustainable Energy Reviews*, May 2016, 58, 1216–1222.
- Décamps, Jean-Paul, Thomas Mariotti, and Stéphane Villeneuve**, “Investment Timing under Incomplete Information,” *Mathematics of Operations Research*, 2005, 30 (2), 472–500.
- Dixit, Robert K. and Robert S. Pindyck**, *Investment under Uncertainty*, Princeton University Press, December 1994.
- DOE**, “FOTW# 1168, Average New Light Truck Price in 2019 Was 43% Higher than the Average New Car Price,” <https://www.energy.gov/eere/vehicles/articles/fotw-1168-january-11-2021-average-new-light-truck-price-2019-was-43-higher> January 2021.

- , “2022 Incremental Purchase Cost Methodology and Results for Clean Vehicles,” December 2022.
- , “FOTW #1327: Annual New Light-Duty EV Sales Topped 1 Million for the First Time in 2023,” <https://www.energy.gov/eere/vehicles/articles/fotw-1327-january-29-2024-annual-new-light-duty-ev-sales-topped-1-million> January 2024.
- , “FOTW #1343: Hybrid Electric Vehicle Sales in the United States Grew 53% in 2023,” <https://www.energy.gov/eere/vehicles/articles/fotw-1343-may-20-2024-hybrid-electric-vehicle-sales-united-states-grew-53> May 2024.
- , “FOTW #1276: U.S. New Light-Duty Vehicle Sales Totaled 13.8 Million in 2022,” <https://www.energy.gov/eere/vehicles/articles/fotw-1276-february-6-2023-us-new-light-duty-vehicle-sales-totaled-138> February 6, 2023.
- DOE**, “Alternative Fuels Data Center: Key Federal Legislation,” [https://afdc.energy.gov/laws/key\\_legislation](https://afdc.energy.gov/laws/key_legislation).
- Fabra, Natalia and Juan-Pablo Montero**, “Technology-Neutral Versus Technology-Specific Procurement,” *The Economic Journal*, February 2023, 133 (650), 669–705.
- Farmer, J. Doyne and François Lafond**, “How Predictable Is Technological Progress?,” *Research Policy*, April 2016, 45 (3), 647–665.
- , **Cameron Hepburn, Penny Mealy, and Alexander Teytelboym**, “A Third Wave in the Economics of Climate Change,” *Climate Research in Gothenburg on*, 2015, 62, 329–357.
- Fedus, William, Carles Gelada, Yoshua Bengio, Marc G. Bellemare, and Hugo Larochelle**, “Hyperbolic Discounting and Learning over Multiple Horizons,” February 2019.

- Funk, Jeffrey L. and Christopher L. Magee**, “Rapid Improvements with No Commercial Production: How Do the Improvements Occur?,” *Research Policy*, April 2015, *44* (3), 777–788.
- Geels, Frank W. and Martina Ayoub**, “A Socio-Technical Transition Perspective on Positive Tipping Points in Climate Change Mitigation: Analysing Seven Interacting Feedback Loops in Offshore Wind and Electric Vehicles Acceleration,” *Technological Forecasting and Social Change*, August 2023, *193*, 122639.
- Gittins, John C.**, “Bandit Processes and Dynamic Allocation Indices,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1979, *41* (2), 148–177.
- , **Richard Weber, and Kevin D. Glazebrook**, *Multi-Armed Bandit Allocation Indices*, second ed., Hoboken, NJ: John Wiley & Sons, 2011.
- Grenadier, Steven R. and Andrey Malenko**, “A Bayesian Approach to Real Options: The Case of Distinguishing between Temporary and Permanent Shocks,” *The Journal of Finance*, 2010, *65* (5), 1949–1986.
- Gunjan, Abhishek and Siddhartha Bhattacharyya**, “A Brief Review of Portfolio Optimization Techniques,” *Artificial Intelligence Review*, May 2023, *56* (5), 3847–3886.
- Gupta, Samarth, Shreyas Chaudhari, Gauri Joshi, and Osman Yağın**, “Multi-Armed Bandits with Correlated Arms,” *IEEE Transactions on Information Theory*, October 2021, *67* (10), 6711–6732.
- Hoehne, Christopher, Matteo Muratori, Paige Jadun, Brian Bush, Arthur Yip, Catherine Ledna, Laura Vimmerstedt, Kara Podkaminer, and Ookie Ma**, “Exploring Decarbonization Pathways for USA Passenger and Freight Mobility,” *Nature Communications*, October 2023, *14* (1), 6913.

- Hou, Chunli, Thomas Little, and Vijay Pant**, “A New Integral Representation of the Early Exercise Boundary for American Put Options,” *Journal of Computational Finance*, March 2000, *3* (3), 73–96.
- Huang, Jing-Zhi, Marti G. Subrahmanyam, and G. George Yu**, “Pricing and Hedging American Options: A Recursive Integration Method,” *The Review of Financial Studies*, January 1996, *9* (1), 277–300.
- Huo, Xiaoguang and Feng Fu**, “Risk-Aware Multi-Armed Bandit Problem with Application to Portfolio Selection,” *Royal Society Open Science*, November 2017, *4* (11), 171377.
- IEA**, “Global EV Outlook 2024,” 2024.
- Ingersoll, Jonathan E.**, *Theory of Financial Decision Making* Rowman & Littlefield Studies in Financial Economics, third (print) ed., Savage, MD: Rowman & Littlefield, 1987.
- Ioannou, Anastasia, Andrew Angus, and Feargal Brennan**, “Risk-Based Methods for Sustainable Energy System Planning: A Review,” *Renewable and Sustainable Energy Reviews*, 2017, *74*, 602–615.
- IPCC**, *Climate Change 2022, Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK: IPCC, 2022.
- Ito, Kiyoshi**, “III. Stochastic Differential Equation and Stochastic Integral Equation,” in “On Stochastic Differential Equations,” Vol. 1, United States: American Mathematical Society, 1951.
- Jun, Tackseung**, “A Survey on the Bandit Problem with Switching Costs,” *De Economist*, December 2004, *152* (4), 513–541.

- Karatzas, Ioannis**, “Gittins Indices in the Dynamic Allocation Problem for Diffusion Processes,” *The Annals of Probability*, February 1984, *12* (1), 173–192.
- Karoui, Nicole EL and Ioannis Karatzas**, “Dynamic Allocation Problems in Continuous Time,” *The Annals of Applied Probability*, May 1994, *4* (2), 255–286.
- Kaspi, Haya and Avi Mandelbaum**, “Lévy Bandits: Multi-Armed Bandits Driven by Lévy Processes,” *The Annals of Applied Probability*, 1995, *5* (2), 541–565.
- Kasy, Maximilian and Alexander Teytelboym**, “Matching with semi-bandits,” *The Econometrics Journal*, January 2023, *26* (1), 45–66.
- Katehakis, Michael N. and Arthur F. Veinott**, “The Multi-Armed Bandit Problem: Decomposition and Computation,” *Mathematics of Operations Research*, 1987, *12* (2), 262–268.
- Keller, Godfrey and Sven Rady**, “Optimal Experimentation in a Changing Environment,” *The Review of Economic Studies*, 1999, *66* (3), 475–507.
- Keppo, Ilkka and Manfred Strubegger**, “Short Term Decisions for Long Term Problems – The Effect of Foresight on Model Based Energy Systems Analysis,” *Energy*, May 2010, *35* (5), 2033–2042.
- Kim, Michael Jong and Andrew E.B. Lim**, “Robust Multiarmed Bandit Problems,” *Management Science*, January 2016, *62* (1), 264–285.
- Kittner, Noah, Ioannis Tsiropoulos, Dalius Tarvydas, Oliver Schmidt, Iain Staffell, and Daniel M. Kammen**, “Chapter 9 - Electric Vehicles,” in Martin Junginger and Atse Louwen, eds., *Technological Learning in the Transition to a Low-Carbon Energy System*, Academic Press, January 2020, pp. 145–163.
- Klein, Manuel**, “Comment on "Investment Timing under Incomplete Information",” *Mathematics of Operations Research*, 2009, *34* (1), 249–254.

- König, Adrian, Lorenzo Nicoletti, Daniel Schröder, Sebastian Wolff, Adam Waclaw, and Markus Lienkamp**, “An Overview of Parameter and Cost for Battery Electric Vehicles,” *World Electric Vehicle Journal*, March 2021, *12* (1), 21.
- Krey, Volker**, “Global Energy-Climate Scenarios and Models: A Review,” *Wiley Interdisciplinary Reviews: Energy and Environment*, 2014, *3* (4), 363–383.
- Kumbaroğlu, Gürkan, Reinhard Madlener, and Mustafa Demirel**, “A Real Options Evaluation Model for the Diffusion Prospects of New Renewable Power Generation Technologies,” *Energy Economics*, July 2008, *30* (4), 1882–1908.
- Lafond, François**, “Forecasting Technological Progress,” March 2025.
- , **Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J. Doyne Farmer**, “How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts,” *Technological Forecasting and Social Change*, March 2018, *128*, 104–117.
- Lai, T.L and Herbert Robbins**, “Asymptotically efficient adaptive allocation rules,” *Advances in Applied Mathematics*, March 1985, *6* (1), 4–22.
- Malekipirbazari, Milad and Özlem Çavuş**, “Index Policy for Multiarmed Bandit Problem with Dynamic Risk Measures,” *European Journal of Operational Research*, January 2024, *312* (2), 627–640.
- Mandelbaum, Avi**, “Continuous Multi-Armed Bandits and Multiparameter Processes,” *The Annals of Probability*, October 1987, *15* (4), 1527–1556.
- Markowitz, Harry**, “Portfolio Selection,” *The Journal of Finance*, 1952, *7* (1), 77–91.
- , “Mean–Variance Approximations to Expected Utility,” *European Journal of Operational Research*, April 2014, *234* (2), 346–355.
- Mazzola, Joseph B. and Kevin F. McCardle**, “A Bayesian Approach to Managing Learning-Curve Uncertainty,” *Management Science*, May 1996, *42* (5), 680.

- McKerracher, Colin**, “China’s Batteries Are Now Cheap Enough to Power Huge Shifts,” *Bloomberg.com*, July 2024.
- McNerney, James, J. Doyne Farmer, and Jessika E. Trancik**, “Historical Costs of Coal-Fired Electricity and Implications for the Future,” *Energy Policy*, June 2011, *39* (6), 3042–3054.
- , – , **Sidney Redner, and Jessika E. Trancik**, “Role of Design Complexity in Technology Improvement,” *Proceedings of the National Academy of Sciences*, May 2011, *108* (22), 9008–9013.
- Mealy, Penny, Pete Barbrook-Johnson, Matthew C. Ives, Sugandha Srivastav, and Cameron Hepburn**, “Sensitive Intervention Points: A Strategic Approach to Climate Action,” *Oxford Review of Economic Policy*, November 2023, *39* (4), 694–710.
- Meng, Jing, Rupert Way, Elena Verdolini, and Laura Diaz Anadon**, “Comparing Expert Elicitation and Model-Based Probabilistic Technology Cost Forecasts for the Energy Transition,” *Proceedings of the National Academy of Sciences of the United States of America*, July 2021, *118* (27).
- Merton, Robert C.**, “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case,” *The Review of Economics and Statistics*, 1969, *51* (3), 247–257.
- Messner, Sabine, Alexander N. Golodnikov, and Andrii Gritsevskii**, “A Stochastic Version of the Dynamic Linear Programming Model MESSAGE III,” *Energy*, September 1996, *21* (9), 775–784.
- Milazzo, Alessandro**, “On the Monotonicity of the Stopping Boundary for Time-Inhomogeneous Optimal Stopping Problems,” *Journal of Optimization Theory and Applications*, October 2024, *203* (1), 336–358.

- Milford, James, Max Henrion, Chad Hunter, Emily Newes, Caroline Hughes, and Samuel F. Baldwin**, “Energy Sector Portfolio Analysis with Uncertainty,” *Applied Energy*, January 2022, *306*, 117926.
- Milovanoff, Alexandre, I. Daniel Posen, and Heather L. MacLean**, “Electrification of Light-Duty Vehicle Fleet Alone Will Not Meet Mitigation Targets,” *Nature Climate Change*, December 2020, *10* (12), 1102–1107.
- Moscarini, Giuseppe and Lones Smith**, “The Optimal Level of Experimentation,” *Econometrica*, 2001, *69* (6), 1629–1644.
- Muratori, Matteo, Paige Jadun, Brian Bush, David Bielen, Laura Vimmerstedt, Jeff Gonder, Chris Gearhart, and Doug Arent**, “Future Integrated Mobility-Energy Systems: A Modeling Perspective,” *Renewable and Sustainable Energy Reviews*, March 2020, *119*, 109541.
- Nagy, Bela, J. Doayne Farmer, Quan M. Bui, and Jessika E. Trancik**, “Statistical Basis for Predicting Technological Progress,” *PLoS ONE*, February 2013, *8* (2), 52669.
- NESO**, “Future Energy Scenarios 2025: Pathways to Net Zero,” 2025.
- Pedersen, Jesper Lund and Goran Peskir**, “On Nonlinear Integral Equations Arising in Problems of Optimal Stopping,” *Proceedings of Functional Analysis*, 2022, *46*, 159–175.
- Peskir, Goran**, “A Change-of-Variable Formula with Local Time on Curves,” *Journal of Theoretical Probability*, July 2005, *18* (3), 499–535.
- **and Albert Shiryaev**, *Optimal Stopping and Free-Boundary Problems* Lectures in Mathematics ETH Zürich, Basel ; Boston: Birkhäuser Verlag, 2006.
- Powell, Warren B.**, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Hoboken, US: John Wiley & Sons, Incorporated, 2011.
- Rodrik, Dani**, “Green Industrial Policy,” *Oxford Review of Economic Policy*, 2014, *30* (3), 469–491.

- Romer, Paul M.**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, 98 (5), S71–S102.
- Salo, Ahti, Michalis Doumpos, Juuso Liesiö, and Constantin Zopounidis**, “Fifty Years of Portfolio Optimization,” *European Journal of Operational Research*, October 2024, 318 (1), 1–18.
- Samuelson, Paul A.**, “Lifetime Portfolio Selection By Dynamic Stochastic Programming,” *The Review of Economics and Statistics*, 1969, 51 (3), 239–246.
- Sani, Amir, Alessandro Lazaric, and Rémi Munos**, “Risk-Aversion in Multi-armed Bandits,” January 2013.
- Schmidt, Oliver, Adam Hawkes, Ajay Gambhir, and Iain Staaell**, “The Future Cost of Electrical Energy Storage Based on Experience Rates,” *Nature Energy*, 2017, 2 (17110).
- Scott, Steven L.**, “A Modern Bayesian Look at the Multi-Armed Bandit,” *Applied Stochastic Models in Business and Industry*, 2010, 26 (6), 639–658.
- Shen, Weiwei, Jun Wang, Yu-Gang Jiang, and Hongyuan Zha**, “Portfolio Choices with Orthogonal Bandit Learning,” *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJAI 2015)*, 2015.
- Sinsel, Simon R., Jochen Markard, and Volker H. Hoffmann**, “How Deployment Policies Affect Innovation in Complementary Technologies—Evidence from the German Energy Transition,” *Technological Forecasting and Social Change*, December 2020, 161.
- Solow, Robert M.**, “Technical Change and the Aggregate Production Function,” *The Review of Economics and Statistics*, 1957, 39 (3), 312–320.
- Sunar, Nur and Zhengli Wang**, “Optimal Multi-armed Bandit with Dependent Arms,” May 2024.

- Unterluggauer, Tim, Jeppe Rich, Peter Bach Andersen, and Seyedmostafa Hashemi**, “Electric Vehicle Charging Infrastructure Planning for Integrated Transportation and Power Distribution Networks: A Review,” *eTransportation*, May 2022, *12*, 100163.
- Way, Rupert, François Lafond, Fabrizio Lillo, Valentyn Panchenko, and J. Doyne Farmer**, “Wright Meets Markowitz: How Standard Portfolio Theory Changes When Assets Are Technologies Following Experience Curves,” *Journal of Economic Dynamics and Control*, April 2019, *101*, 211–238.
- , **Matthew C. Ives, Penny Mealy, and J. Doyne Farmer**, “Empirically Grounded Technology Forecasts and the Energy Transition,” *Joule*, September 2022, *6*, 1–26.
- Weiss, Martin, Andreas Zerfass, and Eckard Helmers**, “Fully Electric and Plug-in Hybrid Cars - An Analysis of Learning Rates, User Costs, and Costs for Mitigating CO2 and Air Pollutant Emissions,” *Journal of Cleaner Production*, March 2019, *212*, 1478–1489.
- Whitehead, Jaramillo P., Suzana Kahn Ribeiro, Peter Newman, Subash Dhar, Ogheneruona E. Diemoudeke, Tsutomu Kajino, David Simon Lee, Sudarmanto Budi Nugroho, Xunmin Ou, Anders Hammer Strømman, and Jake Whitehead**, *Transport*, 1 ed., Cambridge University Press, August
- Whittle, Peter**, “Multi-Armed Bandits and the Gittins Index,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1980, *42* (2), 143–149.
- , “Restless Bandits: Activity Allocation in a Changing World,” *Journal of Applied Probability*, 1988, *25*, 287–298.
- Wright, Theodore Paul**, “Factors Affecting the Cost of Airplanes,” *Journal of the Aeronautical Sciences*, February 1936, *3* (4).

**Yao, Yi-Ching**, "Some Results on the Gittins Index for a Normal Reward Process," *IMS Lecture Notes–Monograph Series Time Series and Related Topics*, January 2006, 52, 284–295.

# Chapter VI

## Conclusion

The three main chapters of this thesis collectively contribute to the dual objective of this thesis “to investigate the reality of future deployment and cost of clean energy technologies while advancing current dynamic models used in the field.” Our findings advance the application of statistical methods in environmental economics and social science while providing insights into the role of technological change in the energy transition.

Throughout the three papers, I underscored the critical role of robust model validation, either on empirical or theoretical grounds. In Chapter III, we used empirical model validation, particularly out-of-sample testing and technology analogues, to identify unreliable findings in the recent literature on the feasibility of future technology cost and deployment. We find that many feasibility assessments are not statistically robust—so far, no one has shown a statistically significant deviation from exponential growth in renewable electricity. In Chapter IV, we demonstrate how similar techniques can be used to develop a robust forecasting model for global and national solar and wind LCOE. For solar, we forecast a continued decline in global LCOE, with large correlations across countries. For wind, we forecast diminishing cost declines due to a stochastic floor cost. In Chapter V, we applied these types of stochastic forecasts in an optimisation setting. Endogenous innovation makes empirical validation difficult. We therefore study this problem theoretically to find the optimal solution under a set of explicit assumptions. We apply the resulting dynamic optimization model to a portfolio of BEVs and ICEs, where BEVs are subject to learning and ICEs are not. For low discount rates, BEVs have crossed a tipping point in 2021, such that it is now optimal to invest all resources into BEVs despite their current higher costs.

This thesis opens the door to further research questions, integrating rigorous dynamic methods with pressing questions in the energy transition. As is often the case, we conclude with more questions than we started with. Below, we summarise key insights for the

energy transition and outline future research directions.

## Key implications for the energy transition

The primary focus of this thesis is the deployment and cost of solar photovoltaics, onshore wind, and battery-electric vehicles. The cost of clean energy technologies is a central determinant of their adoption and implication for the wider economy (Stern, 2014).

In Chapter III, we focused on the deployment of solar photovoltaics and onshore wind and examined recent methods that assess the feasibility of continued rapid growth in these technologies. Through model validation, we identified previous data-based studies on this topic that are not statistically robust. Due to the non-linearity and correlation structures in historical deployment data, empirical feasibility analyses face substantial methodological challenges exacerbated by data limitations. One way of mitigating these limitations is through technology analogies, as we demonstrate with our analysis of mobile phone adoption, internet users, and air travel (see also Farmer and Lafond (2016)). For solar and wind, current deployment data does not exhibit statistically significant signs of a departure from their historically exponential growth trend. Nevertheless, we do expect both technologies to reach an inflection point soon. Particularly for wind, the cost trends uncovered in Chapter IV support the hypothesis that we will see a significant change in growth rates in the next years.

By decomposing renewable electricity costs into their components, Chapter IV develops more precise and granular forecasts for solar and wind LCOE. Our findings suggest a continued strong cost decline for solar across countries, even if a country lags in its deployment. This is not the same for wind. Since balance of system costs have shown no significant decline in the past four decades, wind electricity costs are expected to plateau soon. This is similar to what McNerney et al. (2011) found in the case of coal-fired electricity. While exact floor costs remain uncertain, we do forecast a departure from its current cost trajectory and diminishing returns from the experience curve. Since costs are a significant driver of technology adoption, this result warrants some scepticism around

wind continuing on an exponential growth path without strong political support.

This finding exposes a known contradiction in Wright’s law. [Ferioli et al. \(2009\)](#) previously described the experience curve of a composite technology, made up of multiple other technology components ([Arthur, 2011](#)). If experience rates between components differ substantially, the aggregate and its components cannot all follow Wright’s law. This appears to contradict [Way et al. \(2022\)](#), who argued that floor-cost assumptions have led to biased technology cost scenarios in Integrated Assessment Models. To resolve this contradiction, it is important to highlight three details in Wright’s law.

First is the bias-variance trade-off discussed in Chapter [III](#). The existence of an unknown floor cost in a technology’s future does not justify ad-hoc assumptions about what that floor cost may be. Otherwise, this assumption could introduce bias to the model—a stochastic model without floor costs may still be the preferred predictor. Second, Wright’s law is conditional on future cumulative deployment. While it does not capture all future cost declines ([Lafond et al., 2020](#)), stagnating cumulative deployment still implies reduced cost declines. For wind, this means that if cumulative deployment reaches its asymptote, we may reach the stochastic floor cost without observing it explicitly. Third, the lack of a clear definition of ‘technology’ is a known limitation of Wright’s law ([Ferioli et al., 2009](#)). For two closely related technologies, it may or may not make sense to distinguish them in Wright’s law. The model choice typically depends on the available data, qualitative arguments, and application of the model—as we have seen in Chapter [IV](#). This means that with sufficient innovation in alternative wind energy (e.g., airborne wind energy ([Vermillion et al., 2021](#))), a different learning curve to model the new technology may be appropriate.

Next to wind, BEVs are another technology where a stochastic floor cost can be expected. As discussed in Chapter [V](#), BEVs have exhibited strong long-term growth in their cumulative deployment and a strong decline in their price, primarily due to battery cost experience curves ([Schmidt et al., 2017](#)). However, BEVs also share many components with their ICE counterpart that have not seen a significant cost reduction over the last three decades ([DOE, 2021](#)). This means that BEVs are not only currently more expensive

than ICEs but also that their eventual price difference is limited.

In this context, Chapter V investigates how low BEV costs will have to fall to represent a more favourable investment than ICEs. Based on the Gittins index, our analysis suggests that BEVs already crossed this economic tipping point in 2021. This aligns with empirical findings from the UK secondhand vehicle market (Boulton et al., 2025) and earlier predictions by Schmidt et al. (2017) that anticipated a tipping point at recent battery costs of 150 USD/kWh (Catsaros, 2023). Identifying and anticipating such tipping points is a powerful tool for policymakers, industry, and finance, as it enables targeted interventions that accelerate the energy transition (Mealy et al., 2023).

Before discussing the application of this thesis as a toolkit for decision-makers, it is worth recognising some key limitations of my research findings. In Chapter III, we identified issues with feasibility analyses in the literature but did not model future S-curve deployment of solar and wind. This means that forecasting future technology deployment remains a key gap in the literature that will require additional data collection and statistical rigour. Statistical significance is also a limitation in Chapter IV, where we cannot always make statistically significant model choices. For example, it is an open question if global or regional wind turbine models perform better. This is on top of the existing limitations of Wright's law, such as distinguishing Moore's and Wright's law (Lafond et al., 2020). Additionally, we did not discuss the implications of our forecasts for renewable stakeholders or the root causes of historical and future cost declines. These missing drivers of cost reduction are also a key limitation of the optimisation model developed in Chapter V. We make a number of strong assumptions, such as perfect substitutability, no switching costs, no exogenous innovation, and homogeneous decision-makers that will not always be appropriate. Recognising these limitations is important when considering a future research agenda and applying the results of this thesis in practice.

## New dynamic tools for decision-makers

Through the dual objective of this thesis, we have also demonstrated how robust dynamic models can be applied across research fields. Integrating methods from statistics, probability theory, metrology, and machine learning into socio-economic studies of the energy transition offers novel and reliable insights into future developments.

The two key tools developed in this thesis are the stochastic forecasts for national solar and wind costs in Chapter IV and the dynamic optimization model based on the Gittins index in Chapter V. While there is always room for further refinements to enhance their realism, these models can already serve as valuable aids to decisions in the energy transition. In current times of high uncertainty, economic volatility, and geopolitical instability, the reliability and robustness of dynamic models against short-term fluctuations become even more crucial.

One immediate application of the national renewable cost forecasts from Chapter IV is assessing implications for the renewables value chain. For upstream sectors, such as mining and mineral processing, understanding future technology costs can help set ambitious yet realistic production cost targets and identify new growth opportunities (Fouquet, ed, 2019). This also applies to upstream fossil-fuel sectors facing increasing competition from renewables (Andres et al., 2023). For downstream sectors, such as electricity markets, future technology cost predictions directly influence capacity planning, network expansion, and dispatch strategies (e.g., Unterluggauer et al. (2022)). These applications are inherently linked to policy and regulatory measures that support clean technology deployment and cost reductions. I thus hope that the forecasts presented here will find their way into other tools, such as energy system assessment and integrated assessment models, which are predominantly used to evaluate the interaction between multiple technologies and broader societal objectives.

Policy and system analysis is also a key application of the new dynamic optimization tool developed in Chapter V (Grubb et al., 2024; Barbrook-Johnson et al., 2024). As discussed

in the Chapter, the Gittins index can be used to anticipate economic tipping points and estimate the costs associated with demand-pull policy measures, including the associated uncertainties. As [Mealy et al. \(2024\)](#) highlight, "Invest[ing] in clean energy technologies with evidence of consistent cost declines" is a high-priority sensitive intervention point with the potential for significant societal and environmental impact. While BEVs may have already crossed this threshold (depending on the discount rate), numerous other clean technologies still require policy support. Notable examples include the different clean hydrogen production technologies (e.g., PEM vs AEL electrolyzers), advanced battery chemistries for long-term storage, and synthetic aviation fuels ([Way et al., 2022](#)). Since our method is technology-agnostic, it has broad applicability to these technologies and beyond the energy transition, extending to industrial and innovation strategies more generally.

To maximize the effectiveness of these tools and ensure confidence in model outputs, Chapter [III](#) emphasises the critical importance of rigorous model validation and explicit uncertainty assessments. Without such validation, energy technology research risks propagating misconceptions about technological feasibility, future cost trajectories, and effective investment strategies. In practice, this has frequently led to overly conservative or pessimistic conclusions about the potential for clean energy technology diffusion ([Way et al., 2022](#); [Creutzig et al., 2017](#); [Xiao et al., 2021](#); [Petropoulos et al., 2022](#); [Wagenvoort et al., 2025](#)). Strengthening model validation practices will be essential to improving the accuracy and credibility of future research and applied models in this domain.

## **Future research directions**

Looking ahead, the findings presented in this thesis open multiple avenues for future research, spanning both dynamic modelling and insights into the energy transition.

For the energy transition, a clear research direction is to apply some of our methods and findings to more technologies. Next to solar, wind, and BEVs, many other clean energy technologies will be instrumental to a carbon-neutral energy system, such as hydrogen electrolyzers, grid transmission and batteries, carbon-capture technologies, or nuclear

power. We can apply our LCOE decomposition and forecast methodology in Chapter IV to other technologies as long as there is sufficient national data available. Similarly, our dynamic resource allocation model of Chapter V can be applied to other competitive technologies to identify future techno-economic tipping points.

For solar, wind, BEVs, and other technologies, further exploration of the key drivers behind our forecasts—deployment and cost—as well as their broader implications, remains among the most pressing research priorities. Drivers and implications of technological change are usually closely related, to the point that they can often be indistinguishable (Lafond et al., 2020; Nordhaus, 2014). A primary example thereof is the relationship between a technology’s cost and deployment: Wright’s law predicts cost declines as a function of cumulative deployment, while cost reductions influence deployment itself. Similar relationships hold for other economic and geopolitical factors (Grafström and Poudineh, 2021), such as labour markets, trade networks, and global equity. Many of these factors require further study beyond this thesis. Nevertheless, I hope that this work can build the foundation of future rigorous work in this area.

As highlighted in Chapters III and IV, increasing the regional granularity of technology models can improve our knowledge of the global dynamics driving the energy transition while also making insights more applicable in practice.<sup>1</sup> Future analyses of granular data could consider the role of national policies on clean energy adoption (e.g., Stechemesser et al. (2024)), international knowledge spillovers (e.g., Nemet (2012); Nemet et al. (2020); Gruber (1998); Grafström (2018)), and the impact of national financial markets (e.g., Baer et al. (2022); Egli et al. (2022)). One particularly important research avenue to highlight is the continued study of local tipping points and sensitive intervention points in technology ecosystems (Gallagher et al., 2012). However, identifying actionable interventions will likely require additional context-specific data (Mulugetta et al., 2022).

A second advantage of using granular data for empirical modelling is that it can significantly enhance model development and validation. A major component of this

---

<sup>1</sup>Despite our critique of their methods, Cherp et al. (2021) have contributed significantly to this area by pioneering the use of national data to assess global deployment scenarios.

thesis, particularly in Chapter IV, involved compiling datasets of national clean technology deployment and costs. The value of this dataset extends beyond the models developed here; it can serve as a foundation for future research on national technology dynamics. Combining the granular cost data of this thesis with additional sources will likely yield further important insights for renewables, BEVs, and beyond.

One particularly interesting avenue is that entirely new data types are becoming available due to new methods in machine learning (e.g., use of remote sensing, [Kruitwagen et al. \(2021\)](#)) and AI (e.g., use of large language models, [Cheng et al. \(2024\)](#)). Looking back, some of these methods could have already simplified the work in this thesis and potentially added additional value to the compiled datasets. On top of that, these new datasets allow for more integrated approaches, where the statistical models in focus for this thesis can be combined with more qualitative approaches from expert-based research. A key advantage of AI-based methods in the field is the quantification of qualitative insights.

Machine-learning-based methods also play a central role in advancing statistical models of technological change themselves. These include improving forecasts of technology costs and deployment, expanding the study of Multi-Armed Bandit frameworks and other reinforcement learning techniques, and developing simplified “rules of thumb” to make dynamic models more accessible to policymakers and industry stakeholders.

One fundamental forecasting challenge arises from Chapter III: How can S-curves be used to reliably predict future technology deployment? This is a difficult question due to data limitations, high non-linearity, and the known tendency of S-curve models to produce biased results ([Petropoulos et al., 2022](#)). [Wagenvoort et al. \(2025\)](#) have already made substantial progress on this issue by developing methods to de-bias S-curve forecasts.

Similarly, there is significant potential to refine existing (Moore’s and Wright’s law) forecasting approaches for technology costs. While [Nagy et al. \(2013\)](#) show that existing multi-factor forecasts tend to overfit empirical data, this does not exclude multi-factor models altogether. Instead, it means that more robust multi-factor models require addi-

tional data and more appropriate methods to avoid overfitting (e.g., shrinkage estimators, [Marquardt and Snee \(1975\)](#)). Additionally, future research could integrate technology costs and deployment empirically and further analyse spillover effects between technologies, sectors, and geographical locations.

New forecasting methods can also benefit from integrating novel statistical techniques. Bayesian methods, in particular, are a promising direction that has been underexplored so far (apart from some exceptions, e.g., [Basei et al. \(2024\)](#); [Mazzola and McCardle \(1996\)](#)). These approaches allow model calibration to be informed by explicit prior knowledge derived from expert input, bottom-up modelling, or other data-driven techniques. This brings us back to our methodological choice in the Introduction.

Empirical forecasting results can, in turn, inform portfolio optimization strategies, as demonstrated in Chapter [V](#). Next to the foundational framework applied here, the MAB literature has developed numerous extensions that will be instrumental in integrating new forecasting models and improving the realism of exploration-exploitation trade-offs.

It is worth highlighting a few of these generalisations again since they could be particularly useful in this context. These include incorporating exogenous innovation (such as Moore’s law, potentially in combination with Wright’s law) via the Whittle index ([Whittle, 1988](#)), managing parameter uncertainty through heuristic methods ([Scott, 2010](#); [Li et al., 2023](#)), and accounting for changing demand dynamics and non-linear discounting using composite discount factors ([Fedus et al., 2019](#)). Other generalizations will require the development of new MAB methods, as commonly done in this research domain ([Gittins et al., 2011](#)). Two key questions include how to construct an optimal portfolio of competing and complementary technologies (e.g., solar, wind, and grid storage) and how to determine the ideal number of technologies to support from a broader set of alternatives. Lastly, exploring more diverse and heterogeneous utility structures could enhance the applicability of these methods. For example, incorporating risk aversion into MAB setups is a very active area of reinforcement learning research ([Sani et al., 2013](#); [Malekipirbazari and Çavuş, 2024](#); [Huo and Fu, 2017](#)).

Finally, an important research direction lies in operationalising advanced modelling techniques for use in policy and industry (Lenton et al., 2022). For example, developing ‘rules of thumb’ that reduce complex models into actionable insights is a valuable research contribution. One promising approach in this direction is risk-opportunity analysis, which has emerged as a complement to traditional cost-benefit analyses (Mercure et al., 2021). By explicitly incorporating uncertainty and endogenous dynamics, this approach expands the solution space for decision-makers beyond assessing project costs and benefits that are both difficult to quantify. Instead, it fosters more innovative and risk-tolerant strategies—closely aligned with the risk-seeking behaviors identified in Chapter V (see also Rodrik (2014)). This research area is closely linked to the effective communication of findings. Bridging the gap between complex models, practical decision-making, and the public can enable more strategic choices in industrial, trade, and fiscal policy (Barbrook-Johnson et al., 2024).

## Concluding remarks

This thesis was written at a time when the world stands at a pivotal crossroads—a rapid energy transition is becoming increasingly feasible and economically advantageous (Way et al., 2022) yet achieving global climate targets remains an immense challenge (Bevacqua et al., 2025; Cannon, 2025). Over the past decade, we have crossed multiple economic tipping points, with solar panels, wind turbines, and batteries reshaping the economics of the energy sector and creating unprecedented opportunities for decarbonization at scale.

Seizing this opportunity for transformative global change demands rigorous, evidence-based decision-making, grounded in trusted models and a forward-looking perspective. By integrating the key insights from this research with ongoing innovations and by proactively addressing the open questions through future work, policymakers and scholars can collaboratively guide the global energy system toward a more sustainable, affordable, and secure future. The energy transition, as demonstrated in this study, is not a distant goal—it is a rapidly unfolding reality, one that can be accelerated through informed action and continued curiosity.

## References

- Andres, Pia, Penny Mealy, Nils Handler, and Samuel Fankhauser**, “Stranded Nations? Transition Risks and Opportunities towards a Clean Economy,” *Environmental Research Letters*, April 2023, 18 (4), 045004.
- Arthur, W. Brian**, *The Nature of Technology: What It Is and How It Evolves*, first trade paperback ed., New York, London, Toronto, Sydney: Free Press, 2011.
- Baer, Moritz, Ben Caldecott, Jacob Kastl, Alissa M. Kleinnijenhuis, and Nicola Ranger**, “TRISK - A Climate Stress Test for Transition Risk,” October 2022.
- Barbrook-Johnson, Pete, Jean-François Mercure, Simon Sharpe, Cristina Peñasco, Cameron Hepburn, Laura Diaz Anadon, J. Doyne Farmer, and Timothy M. Lenton**, “Economic Modelling Fit for the Demands of Energy Decision Makers,” *Nature Energy*, March 2024, 9 (3), 229–231.
- Basei, Matteo, Giorgio Ferrari, and Neofytos Rodosthenous**, “Uncertainty over Uncertainty in Environmental Policy Adoption: Bayesian Learning of Unpredictable Socioeconomic Costs,” *Journal of Economic Dynamics and Control*, April 2024, 161, 104841.
- Bevacqua, Emanuele, Carl-Friedrich Schleussner, and Jakob Zscheischler**, “A Year above 1.5 °C Signals That Earth Is Most Probably within the 20-Year Period That Will Reach the Paris Agreement Limit,” *Nature Climate Change*, March 2025, 15 (3), 262–265.
- Boulton, Chris A., Joshua E. Buxton, and Timothy M. Lenton**, “Early Opportunity Signals of a Tipping Point in the UK’s Second-Hand Electric Vehicle Market,” *Earth System Dynamics*, March 2025, 16 (2), 411–421.
- Cannon, Alex J.**, “Twelve Months at 1.5 °C Signals Earlier than Expected Breach of Paris Agreement Threshold,” *Nature Climate Change*, March 2025, 15 (3), 266–269.

**Catsaros, Oktavia**, “Lithium-Ion Battery Pack Prices Hit Record Low of \$139/kWh,” November 2023.

**Cheng, Zhi-Qi, Yifei Dong, Aike Shi, Wei Liu, Yuzhi Hu, Jason O’Connor, Alexander G. Hauptmann, and Kate S. Whitefoot**, “SHIELD: LLM-Driven Schema Induction for Predictive Analytics in EV Battery Supply Chain Disruptions,” October 2024.

**Cherp, Aleh, Vadim Vinichenko, Jale Tosun, Joel A. Gordon, and Jessica Jewell**, “National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets,” *Nature Energy*, 2021, 6, 742–754.

**Creutzig, Felix, Peter Agoston, Jan Christoph Goldschmidt, Gunnar Luderer, Gregory F. Nemet, and Robert C. Pietzcker**, “The Underestimated Potential of Solar Energy to Mitigate Climate Change,” *Nature Energy*, August 2017, 2 (9).

**DOE**, “FOTW# 1168, Average New Light Truck Price in 2019 Was 43% Higher than the Average New Car Price,” <https://www.energy.gov/eere/vehicles/articles/fotw-1168-january-11-2021-average-new-light-truck-price-2019-was-43-higher> January 2021.

**Egli, Florian, Friedemann Polzin, Mark Sanders, Tobias Schmidt, Alexandra Serebriakova, and Bjarne Steffen**, “Financing the Energy Transition: Four Insights and Avenues for Future Research,” *Environmental Research Letters*, May 2022, 17 (5), 051003.

**Farmer, J. Doyne and François Lafond**, “How Predictable Is Technological Progress?,” *Research Policy*, April 2016, 45 (3), 647–665.

**Fedus, William, Carles Gelada, Yoshua Bengio, Marc G. Bellemare, and Hugo Larochelle**, “Hyperbolic Discounting and Learning over Multiple Horizons,” February 2019.

- Ferioli, Francesco, Koen Schoots, and Bob van der Zwaan**, “Use and Limitations of Learning Curves for Energy Technology Policy: A Component-Learning Hypothesis,” *Energy Policy*, July 2009, *37* (7), 2525–2535.
- Fouquet, Roger, ed.**, *Handbook on Green Growth*, Cheltenham, UK ; Northampton, MA: Edward Elgar Publishing, 2019.
- Gallagher, Kelly Sims, Arnulf Grübler, Laura Kuhl, Gregory F. Nemet, and Charlie Wilson**, “The Energy Technology Innovation System,” *Annual Review of Environment and Resources*, November 2012, *37* (1), 137–162.
- Gittins, John C., Richard Weber, and Kevin D. Glazebrook**, *Multi-Armed Bandit Allocation Indices*, second ed., Hoboken, NJ: John Wiley & Sons, 2011.
- Grafström, Jonas**, “International Knowledge Spillovers in the Wind Power Industry: Evidence from the European Union,” *Economics of Innovation and New Technology*, April 2018, *27* (3), 205–224.
- **and Rahmatallah Poudineh**, “A Critical Assessment of Learning Curves for Solar and Wind Power Technologies,” Technical Report 43, The Oxford Institute for Energy Studies, Oxford 2021.
- Grubb, Michael, Rutger-Jan Lange, Nicolas Cerkez, Ida Sognaes, Claudia Wieners, and Pablo Salas**, “Dynamic Determinants of Optimal Global Climate Policy,” *Structural Change and Economic Dynamics*, December 2024, *71*, 490–508.
- Gruber, Harald**, “Learning by Doing and Spillovers: Further Evidence for the Semiconductor Industry,” *Review of Industrial Organization*, 1998, *13*, 697–711.
- Huo, Xiaoguang and Feng Fu**, “Risk-Aware Multi-Armed Bandit Problem with Application to Portfolio Selection,” *Royal Society Open Science*, November 2017, *4* (11), 171377.

- Kruitwagen, L., K. T. Story, J. Friedrich, L. Byers, S. Skillman, and C. Hepburn,** “A Global Inventory of Photovoltaic Solar Energy Generating Units,” *Nature*, October 2021, *598* (7882), 604–610.
- Lafond, François, Diana Seave Greenwald, and J. Doyne Farmer,** “Can Stimulating Demand Drive Costs Down? World War II as a Natural Experiment,” *SSRN Electronic Journal*, June 2020.
- Lenton, Timothy M., Scarlett Benson, Talia Smith, Theodora Ewer, Victor Lanel, Elizabeth Petykowski, Thomas W. R. Powell, Jesse F. Abrams, Fenna Blomsma, and Simon Sharpe,** “Operationalising Positive Tipping Points towards Global Sustainability,” *Global Sustainability*, January 2022, *5*, e1.
- Li, Xiao, Yuqiang Li, and Xianyi Wu,** “Empirical Gittins Index Strategies with  $\varepsilon$ -Explorations for Multi-Armed Bandit Problems,” *Computational Statistics & Data Analysis*, April 2023, *180*, 107610.
- Malekipirbazari, Milad and Özlem Çavuş,** “Index policy for multiarmed bandit problem with dynamic risk measures,” *European Journal of Operational Research*, January 2024, *312* (2), 627–640.
- Marquardt, Donald W. and Ronald D. Snee,** “Ridge Regression in Practice,” *The American Statistician*, February 1975, *29* (1), 3–20.
- Mazzola, Joseph B. and Kevin F. McCardle,** “A Bayesian Approach to Managing Learning-Curve Uncertainty,” *Management Science*, May 1996, *42* (5), 680.
- McNerney, James, J. Doyne Farmer, and Jessika E. Trancik,** “Historical Costs of Coal-Fired Electricity and Implications for the Future,” *Energy Policy*, June 2011, *39* (6), 3042–3054.
- Mealy, Penny, Michael Ganslmeier, Catrina Godinho, and Stephane Hallegatte,** “Empirical Identification of Feasible and Strategic Climate Policies,” January 2024.

- , Pete Barbrook-Johnson, Matthew C. Ives, Sugandha Srivastav, and Cameron Hepburn, “Sensitive Intervention Points: A Strategic Approach to Climate Action,” *Oxford Review of Economic Policy*, November 2023, 39 (4), 694–710.
- Mercure, Jean-Francois, Simon Sharpe, Jorge E. Vinuales, Matthew Ives, Michael Grubb, Aileen Lam, Paul Drummond, Hector Pollitt, Florian Knobloch, and Femke J. M. M. Nijse, “Risk-Opportunity Analysis for Transformative Policy Design and Appraisal,” *Global Environmental Change*, September 2021, 70, 102359.
- Mulugetta, Yacob, Youba Sokona, Philipp A. Trotter, Samuel Fankhauser, Jessica Omukuti, Lucas Somavilla Croxatto, Bjarne Steffen, Meron Tesfamichael, Edo Abraham, Jean-Paul Adam, Lawrence Agbemabiese, Churchill Agutu, Mekalia Paulos Aklilu, Olakunle Alao, Bothwell Batidzirai, Getachew Bekele, Anteneh G. Dagnachew, Ogunlade Davidson, Fatima Denton, E. Ogheneruona Diemuodeke, Florian Egli, Gebrekidan Gebresilassie Eshetu, Muluaem Gebreslassie, Mamadou Goundiam, Haruna Kachalla Gujba, Yohannes Hailu, Adam D. Hawkes, Stephanie Hirmer, Helen Hoka, Mark Howells, Abdulrasheed Isah, Daniel Kammen, Francis Kemausuor, Ismail Khennas, Wikus Kruger, Ifeoma Malo, Linus Mofor, Minette Nago, Destenie Nock, Chukwumerije Okereke, S. Nadia Ouedraogo, Benedict Probst, Maria Schmidt, Tobias S. Schmidt, Carlos Shenga, Mohamed Sokona, Jan Christoph Steckel, Sebastian Sterl, Bernard Tembo, Julia Tomei, Peter Twesigye, Jim Watson, Harald Winkler, and Abdulmutalib Yussuff, “Africa Needs Context-Relevant Evidence to Shape Its Clean Energy Future,” *Nature Energy*, October 2022, 7, 1015–1022.
- Nagy, Bela, J. Doyne Farmer, Quan M. Bui, and Jessika E. Trancik, “Statistical Basis for Predicting Technological Progress,” *PLoS ONE*, February 2013, 8 (2), 52669.
- Nemet, Gregory F., “Inter-Technology Knowledge Spillovers for Energy Technologies,” *Energy Economics*, September 2012, 34 (5), 1259–1270.

– , **Jiaqi Lu, Varun Rai, and Rohan Rao**, “Knowledge Spillovers between PV Installers Can Reduce the Cost of Installing Solar PV,” *Energy Policy*, September 2020, *144*, 111600.

**Nordhaus, William D.**, “The Perils of the Learning Model for Modeling Endogenous Technological Change,” *Energy Journal*, January 2014, *35* (1), 1–13.

**Petropoulos, Fotios, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K. Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J. Bessa, Jakub Bijak, John E. Boylan, Jethro Browell, Claudio Carnevale, Jennifer L. Castle, Pasquale Cirillo, Michael P. Clements, Clara Cordeiro, Fernando Luiz Cyrino Oliveira, Shari De Baets, Alexander Dokumentov, Joanne Ellison, Piotr Fiszeder, Philip Hans Franses, David T. Frazier, Michael Gilliland, M. Sinan Gönül, Paul Goodwin, Luigi Grossi, Yael Grushka-Cockayne, Mariangela Guidolin, Massimo Guidolin, Ulrich Gunter, Xiaojia Guo, Renato Guseo, Nigel Harvey, David F. Hendry, Ross Hollyman, Tim Januschowski, Jooyoung Jeon, Victor Richmond R. Jose, Yanfei Kang, Anne B. Koehler, Stephan Kolassa, Nikolaos Kourentzes, Sonia Leva, Feng Li, Konstantia Litsiou, Spyros Makridakis, Gael M. Martin, Andrew B. Martinez, Sheik Meeran, Theodore Modis, Konstantinos Nikolopoulos, Dilek Önköl, Alessia Paccagnini, Anastasios Panagiotelis, Ioannis Panapakidis, Jose M. Pavía, Manuela Pedio, Diego J. Pedregal, Pierre Pinson, Patrícia Ramos, David E. Rapach, J. James Reade, Bahman Rostami-Tabar, Michał Rubaszek, Georgios Sermpinis, Han Lin Shang, Evangelos Spiliotis, Aris A. Syntetos, Priyanga Dilini Talagala, Thiyanga S. Talagala, Len Tashman, Dimitrios Thomakos, Thordis Thorarinsdottir, Ezio Todini, Juan Ramón Traperó Arenas, Xiaoqian Wang, Robert L. Winkler, Alisa Yusupova, and Florian Ziel**, “Forecasting: Theory and Practice,” *International Journal of Forecasting*, July 2022, *38* (3), 705–871.

- Rodrik, Dani**, “Green Industrial Policy,” *Oxford Review of Economic Policy*, 2014, 30 (3), 469–491.
- Sani, Amir, Alessandro Lazaric, and Rémi Munos**, “Risk-Aversion in Multi-armed Bandits,” January 2013.
- Schmidt, Oliver, Adam Hawkes, Ajay Gambhir, and Iain Staaell**, “The Future Cost of Electrical Energy Storage Based on Experience Rates,” *Nature Energy*, 2017, 2 (17110).
- Scott, Steven L.**, “A Modern Bayesian Look at the Multi-Armed Bandit,” *Applied Stochastic Models in Business and Industry*, 2010, 26 (6), 639–658.
- Stechemesser, Annika, Nicolas Koch, Ebba Mark, Elina Dilger, Patrick Klösel, Laura Menicacci, Daniel Nachtigall, Felix Pretis, Nolan Ritter, Moritz Schwarz, Helena Vossen, and Anna Wenzel**, “Climate Policies That Achieved Major Emission Reductions: Global Evidence from Two Decades,” *Science*, August 2024, 385 (6711), 884–892.
- Stern, Nicholas H.**, *The Economics of Climate Change: The Stern Review*, Cambridge: Cambridge University Press, 2014.
- Unterluggauer, Tim, Jeppe Rich, Peter Bach Andersen, and Seyedmostafa Hashemi**, “Electric Vehicle Charging Infrastructure Planning for Integrated Transportation and Power Distribution Networks: A Review,” *eTransportation*, May 2022, 12, 100163.
- Vermillion, Chris, Mitchell Cobb, Lorenzo Fagiano, Rachel Leuthold, Moritz Diehl, Roy S. Smith, Tony A. Wood, Sebastian Rapp, Roland Schmehl, David Olinger, and Michael Demetriou**, “Electricity in the Air: Insights from Two Decades of Advanced Control Research and Experimental Flight Testing of Airborne Wind Energy Systems,” *Annual Reviews in Control*, January 2021, 52, 330–357.

**Wagenvoort, Benjamin, Joel Dyer, François Lafond, and J. Doyne Farmer**, “The Universality of Technology Production,” 2025.

**Way, Rupert, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer**, “Empirically Grounded Technology Forecasts and the Energy Transition,” *Joule*, September 2022, *6*, 1–26.

**Whittle, Peter**, “Restless Bandits: Activity Allocation in a Changing World,” *Journal of Applied Probability*, 1988, *25*, 287–298.

**Xiao, Mengzhu, Tobias Junne, Jannik Haas, and Martin Klein**, “Plummeting Costs of Renewables - Are Energy Scenarios Lagging?,” *Energy Strategy Reviews*, May 2021, *35*, 100636.

# Appendix A

## Supplementary Materials for *The need for better statistical testing in data-driven energy technology modeling*

### A.1 Supplementary discussion on the examined statistical model selected from recent literature

#### A.1.1 Detailed explanation of the model

For better understanding of the points raised in the main text and following subsections, we present here a summary of the methodology used by [Cherp et al. \(2021\)](#). Their procedure is to:

1. Select the 60 countries that currently have the largest electricity systems. (The purpose of this is to eliminate countries with high volatility due to a small number of projects, but also countries with significant electricity import or export; this excludes countries such as Luxembourg and Morocco).
2. For each country, exclude deployment data points if they are less than 0.1% of the electricity supply.
3. Exclude countries with fewer than 6 data points remaining.
4. Exclude countries with fewer than 2 deployment data points less than 1% of the electricity supply. (Cherp et al. define the first year for which wind/solar electricity production exceeds 1% of the total electricity supply as the take-off year.)
5. Fit a logistic S-curve  $S(t)$  to the deployment data of the remaining countries, i.e., those that have not been eliminated by the preceding steps.

6. Let  $f(t)$  be the percentage of deployment at time  $t$  relative to the asymptote of the fitted logistic S-curve (i.e.  $f(t) = S(t)/L$ ). Let the *maturity* be  $m = f(t_{max})$ , where  $t_{max}$  is the maximum time of the analysis. Exclude any countries where  $m < 0.5$ . (The proposed purpose is to eliminate any countries whose S-curve fit indicates that the maximum growth rate lies in the future.) Note that the maturity is based on the logistic fit, not on the actual data – the two can be very different.
7. For each country, calculate  $t_{10}$ , defined as the time when  $f(t_{10}) = 0.1$ , and  $t_{90}$ , defined as  $f(t_{90}) = 0.9$  (allowing cases where  $t_{90} > t_{max}$ ). For wind, exclude countries where  $t_{90} - t_{10} < 4$  years, and for solar, exclude countries where  $t_{90} - t_{10} < 2$  years. (The proposed purpose is to eliminate any countries whose S-curve fits show short spikes, indicating “one-off growth spurts”. This applies to 6 countries in 2019.)
8. Fit a Gompertz S-curve to the remaining countries.
9. For each remaining country, calculate the maximum growth rate of the fitted logistic S-curve and the maximum growth rate of the fitted Gompertz S-curve. The average of these two numbers is the estimated historic maximum growth rate  $G$  used by Cherp et al.

As evident from this description, the sampling procedure based on the maturity  $m$  and times  $t_{10}$ ,  $t_{90}$  depends solely on the fitted logistic S-curve. It does not include any identification of statistical significance.

### A.1.2 The reliability of inflection point estimates

In testing the Cherp et al. method using their data, we find that the estimated inflection point  $t_0$  (i.e., the year of maximum growth) is frequently increasing for both wind and solar – shown in Figure A.1. For the countries that survive their filtering process,  $t_0$  values are highly non-stationary, meaning that the mean values change systematically with time. For their method to be reliable, when we increment  $t_{max}$ , the  $t_0$  values of countries that

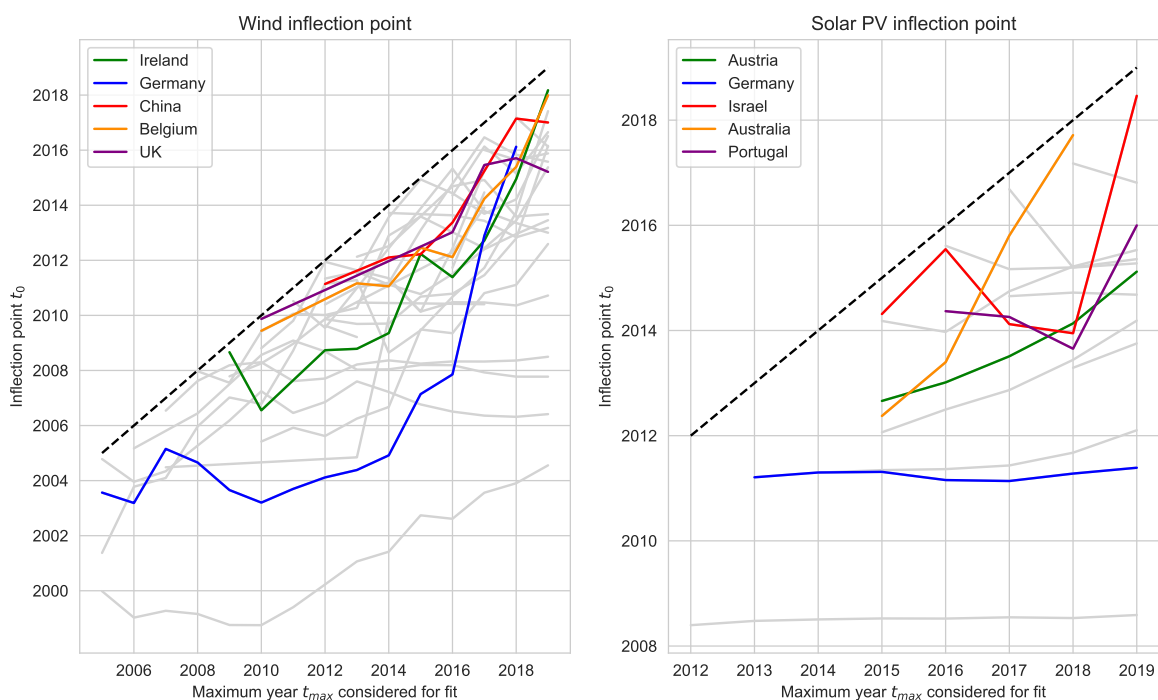


Figure A.1: **Inflection points  $t_0$  as calculated by a logistic fit used in Cherp et al. (2021) versus the maximum year considered for the fit ( $t_{max}$ ).** Each line represents the inflection point  $t_0$  calculated for one country at different values of  $t_{max}$  based on a logistic S-curve fit and filtered according to the method described in the main text. For comparison, the identity line is shown in dashed black. There are, by definition, no values above the identity line since the method of Cherp et al. discards countries for which  $t_0$  lies in the future. The tendency of  $t_0$  to increase with  $t_{max}$  is primarily driven by changes due to countries already in the sample rather than new entries.

are already in the sample should remain roughly constant, and most of the larger values of  $t_0$  should come from new countries entering the sample. This is not what we observe. Instead, in most cases the value of  $t_0$  for countries that are already in the sample keeps moving into the future as  $t_{max}$  increases. For example, for wind power in Germany,  $t_0$  was estimated to be 2004 for an analysis in 2005, but estimated to be about 2016 for an analysis in 2018 – and then Germany leaves the sample in 2019 because  $t_0$  is estimated to be even further in the future. There are only a few countries where  $t_0$  remains constant, and cases where it decreases over time are rare. The non-stationarity of the estimated value of  $t_0$  is illustrated by the fact that when  $t_{max}$  increases by a year, the average value of  $t_0$  increases by 0.6 years for wind and 0.3 years for solar. Across the length of the sample, the estimated national inflection point increased by 8.4 years for wind and 2.1 years for solar.

### A.1.3 The application of the examined statistical model to other technologies

As explained in the main text, the fact that the method of [Cherp et al. \(2021\)](#) is technology-agnostic means that we can test its reliability on other technologies. In this section, we provide additional details to the mobile phone data analysis described in the main text. To ensure that our analysis does not rely on one-off particularities of the mobile phone market and technology, we have also tested the method of [Cherp et al. \(2021\)](#) on national data for annual aviation passenger kilometres and internet users. The results are shown in [Figure A.3](#).

For all three technologies, in all countries, deployment is close or at their S-curve saturation, shown in the bottom row of [Figure A.3](#). The same is true for the global average. We can, therefore, set our threshold for the technology take-off with respect to the maximum saturation reached in the time series. For simplification, we set the same threshold for all technologies. The procedure is as follows:

1. For each technology, let  $\hat{L}$  be the maximum global deployment (normalized).
2. For each technology, let the take-off threshold in each country be  $\hat{L} * 2.5\%$ . (This is equivalent to the 1% threshold for wind and solar PV.)
3. For each technology, exclude data prior to the formative threshold in each country, defined as  $\hat{L} * 0.25\%$ . (This is equivalent to the 0.1% threshold for wind and solar PV.)

We have also performed a sensitivity analysis around these thresholds and do not find a significant difference in the results

#### **Additional information for the mobile phone analysis**

As stated in the main text, the mobile phone analysis uses national deployment data of mobile cellular telephone subscriptions provided by the International Telecommunication

Union [ITU \(2022\)](#). We normalize the number of phone subscriptions to the national population (or to the global population when considering the global data). The mobile phone dataset includes the national number of active contracts that provide access to the Public Switched Telephone Network (PSTN) using cellular technology. This includes all post-paid and active pre-paid subscriptions that offer voice communication, excluding mobile-data-only subscriptions.

Figure [A.2](#) shows logistic S-curves fitted to the global data for mobile phones. The fit yields a maximum growth rate  $G$  of roughly 9% at the inflection point in 2008 and an adoption asymptote  $L$  of about 123% of the total population, indicating that many people have more than one subscription. The adoption levels of solar and wind are by definition bounded at less than 100% of total electricity generation, so given the difference in normalization, we should expect *a priori* that the maximum growth rates  $G = kL/4$  for solar and wind will be smaller than those for mobile phones. However, our purpose is to test the performance of the Cherp et al. method at forecasting growth rates, which should not depend on the asymptotic level. The calculated exponential growth rate for mobile phones  $k$  is about 30%, which is comparable to the compound annual growth rates for solar and wind energy that we quoted above. (See also ([IEA, 2022a,b](#))).

### Aviation data analysis

The aviation analysis is based on the CHAT dataset ([Comin et al., 2006](#)). We consider 54 countries between 1930 and 1993 and normalize by the population, as in the mobile phone example. As shown in the bottom-left panel of Figure [A.3](#), the passenger kilometers traveled are largely saturated on global and most national levels, with the global S-curve inflection point estimated to be 1985 (based on a logistic growth curve fit). The national growth rate sample is selected using the previously discussed procedure, including countries with at least 6 data points above 8 km per person and 2 data points above 8km per person.

One caveat of this dataset is the fact that national passenger-km is calculated based on the country of registration of the aviation company. This may be different than

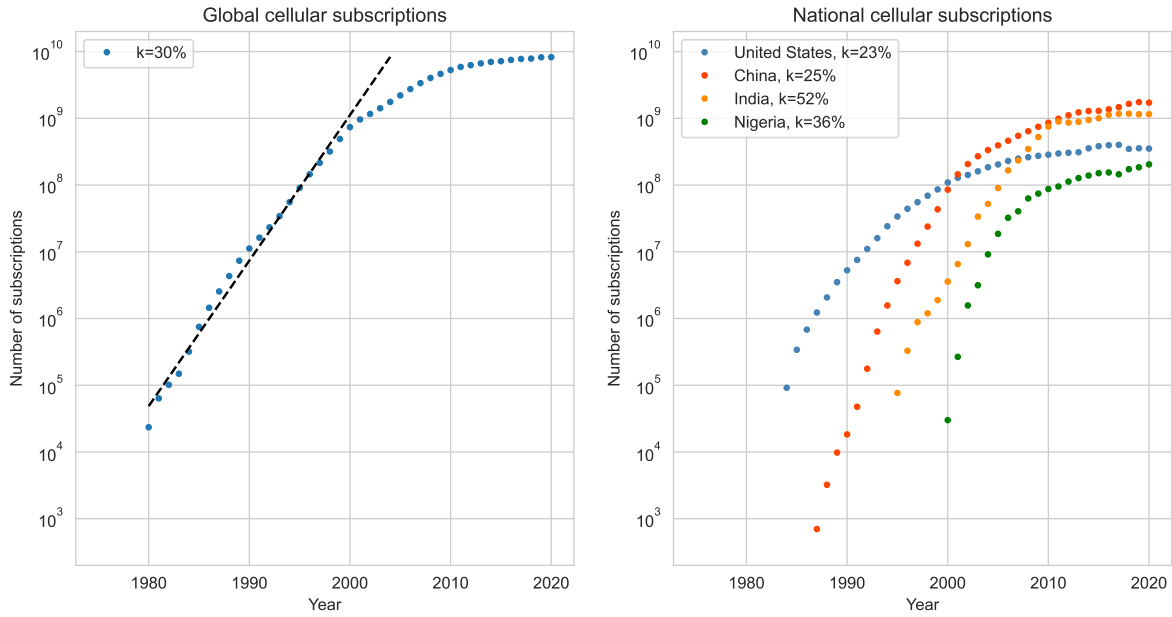


Figure A.2: **Global (left) and selected national (right) deployment data for mobile cellular subscriptions.** Both global and national diffusion data show stagnating numbers of subscriptions, indicating that the S-curved diffusion is mostly complete. Data are shown in logarithmic scale.  $k$  denotes the exponential growth rate estimated through a logistic S-curve fit. The historic global deployment data until 2000 shows approximate exponential growth; the dashed black line corresponding to exponential growth (based on linear regression on a subset of the data up until 2000) is shown for comparison.

the nationality of the passengers, which we normalize by. While this difference blurs the quantitative results of our analysis somewhat, the non-stationarity of the resulting distribution of maximum growth rates  $G$  still indicates non-stationarity in the diffusion process (shown in the upper-left panel of Figure A.3).

Between 1970 and 1980, the median value of  $G$  jumps from around 5 to 35 km per person. Only after the global inflection point does the distribution stabilize at around 15 km per person. This is consistent with the actual global maximum growth rate of 16 km per person. We can also see that the analysis using the [Cherp et al. \(2021\)](#) method is biased towards lower values of the maximum historic growth rate  $G$ , particularly for early years.

## Internet user data analysis

The internet user analysis is based on data provided by the [ITU \(2022\)](#). We consider 148 countries between 1990 and 2021. As shown in the bottom-middle panel of Figure

[A.3](#), the number of internet users is largely saturated on global and most national levels, with the global S-curve inflection point estimated to be 2012 (based on a logistic growth curve fit). The national growth rate sample is selected using an asymptote of  $L = 100\%$ , including countries with at least 6 data points above 0.1% and 2 data points above 2.5% penetration.

As before, we see large non-stationarity in the maximum historic growth rate  $G$  sample (shown in the upper-middle panel of [Figure A.3](#)). In 1995, the median maximum growth rate was only around 1%, growing to over 5% after the global inflection point. This is consistent with the actual global maximum growth rate of 3%. The downward bias is less visible in the distribution median but becomes apparent when we consider the interquartile range.

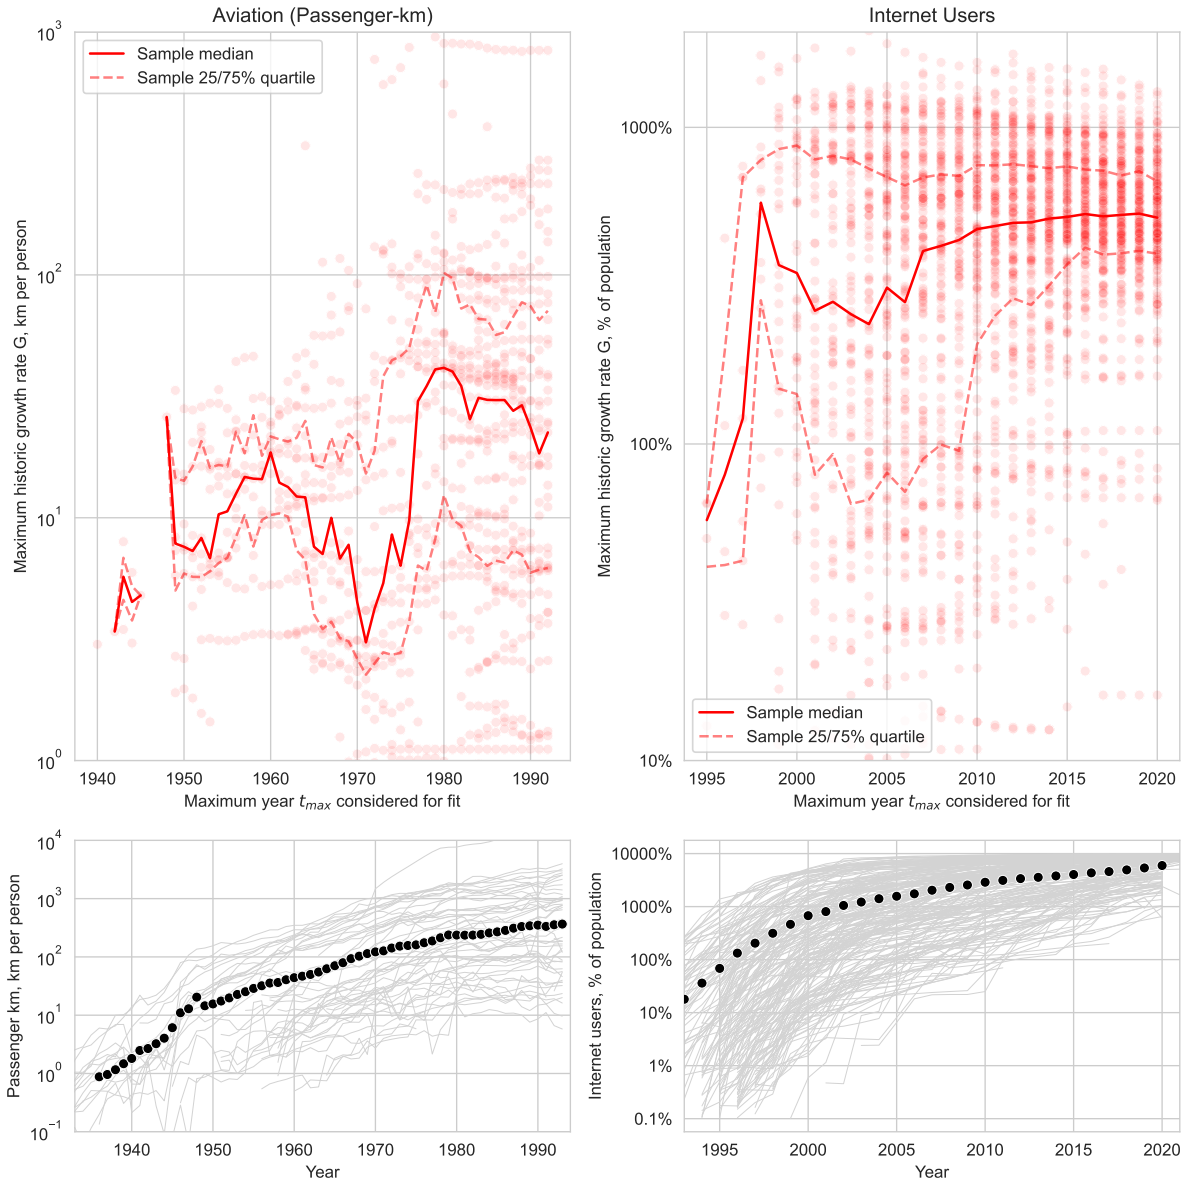


Figure A.3: **Distribution of national maximum growth rates ( $G$ ) for airline passenger km and internet users vs. the maximum year considered for the S-curve fit ( $t_{max}$ ).** (Top panels) National maximum growth rates included in the sample according to the filtering applied by [Cherp et al. \(2021\)](#) are shown as red dots. The red line shows the median maximum growth rate in the sample. Dashed lines represent the interquartile range of the respective distributions. (Bottom panels) National deployment of each technology, normalized by national population, is shown in grey. The black dots show the global deployment, normalized by the global population. Both technologies show sub-linear behavior close to global and most national saturation.

## A.2 Additional Figures

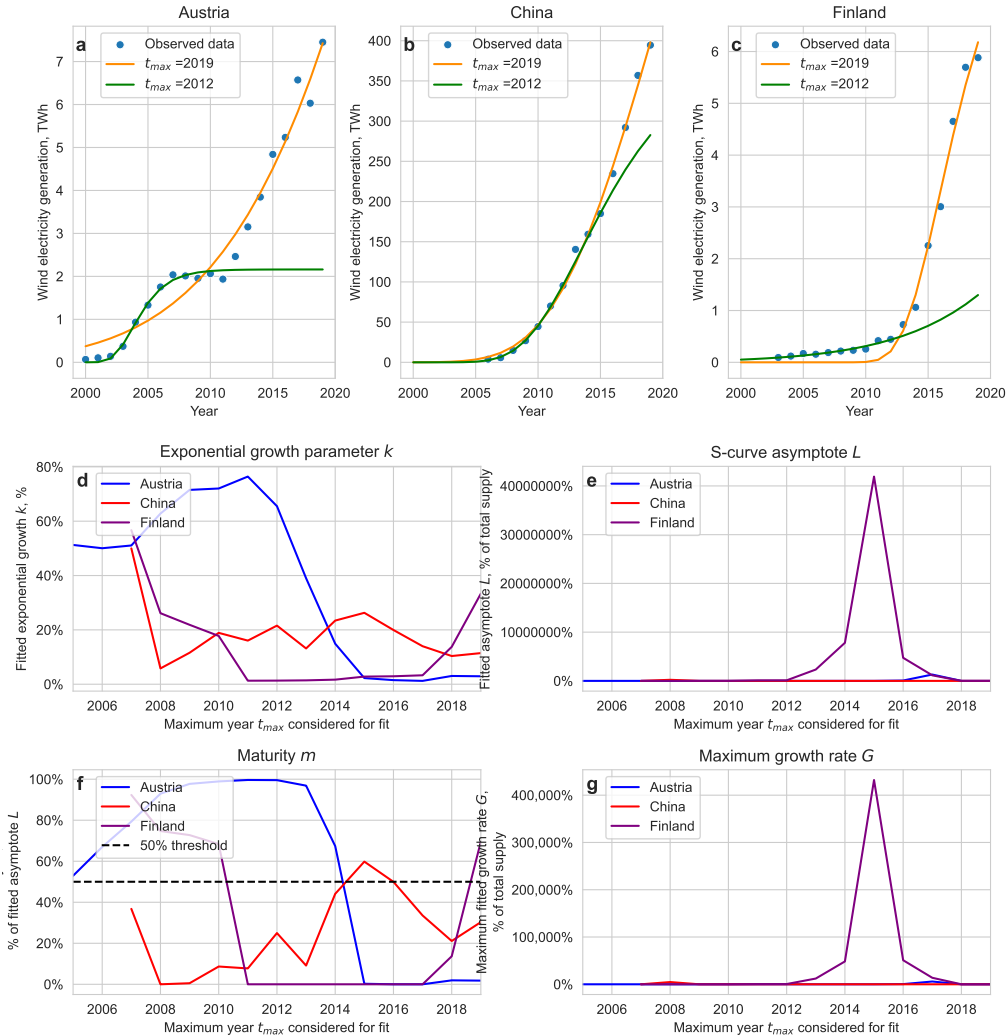


Figure A.4: **S-curve fits for electricity generation from (onshore) wind in Austria, China, and Finland for different maximum years  $t_{max}$ .** (a-c) The deployment of onshore wind in (a) Austria, (b) China, and (c) Finland shows intervals of rapid acceleration and deceleration. With the exception of China, the Gompertz S-curve fit based on data up until  $t_{max} = 2012$  (green) differs starkly from that including data up until  $t_{max} = 2019$  (orange). (d) The exponential growth parameter  $k$  of the resulting S-curve fit varies widely. (e) Particularly for the S-curve asymptote  $L$ , we observe drastic changes depending on the length of the time series. (f) The maturity  $m$  as a percentage of the asymptote  $L$  for the Gompertz fit (not used in the method of [Cherp et al. \(2021\)](#)) (g) The maximum growth rate  $G = kL/4$  reflects the instability in the S-curve parameters  $k$  and  $L$ .

## References

- Cherp, Aleh, Vadim Vinichenko, Jale Tosun, Joel A Gordon, and Jessica Jewell**, “National Growth Dynamics of Wind and Solar Power Compared to the Growth Required for Global Climate Targets,” *Nature Energy*, 2021, 6, 742–754.
- Comin, Diego, Bart Hobijn, and Emilie Rovito**, “Five Facts You Need to Know About Technology Diffusion,” January 2006.
- IEA**, “Solar PV,” <https://www.iea.org/reports/solar-pv> September 2022.
- , “Wind Electricity,” <https://www.iea.org/reports/wind-electricity> September 2022.
- ITU**, “ITU DataHub,” <https://datahub.itu.int/data/?i=178> 2022.

# Appendix B

## Supplementary Materials for *Will national renewable costs continue declining?*

### B.1 Notation overview

Variable	Name	Description	Unit
$\mathcal{L}$	Levelized Cost of Electricity (LCOE)	Average annual cost to produce 1 kWh of electricity	$USD/kWh$
$\mathcal{I}$	Initial investment cost	Initial overnight investment cost of a new wind/solar PV asset	$USD/kW$
$O$	Operating & Maintenance cost (O&M)	Expected annual O&M cost of a new wind/solar PV asset, averaged over the asset lifetime	$USD/kW/year$
$\mathcal{C}$	Capital Recovery Factor	Annualized return required to meet the cost of debt and equity for the initial investment $\mathcal{I}$ , averaged over the asset lifetime (present value)	$\%/year$
$\mathcal{Y}$	Capacity Factor	Annual yield of a wind/solar PV asset	$kWh/kW/year$
$\mathcal{D}$	Weighted average Cost of Capital (WACC)	Annual return on the initial investment $\mathcal{I}$ , averaged over investors, weighted by the share of investment	$\%/year$
$T$	Asset lifetime	Financial & operating lifetime of the wind/solar PV asset	years
$Y$	Electricity generation	Electricity generated within a given year	$kWh/year$
$\bar{Y}$	Average annual electricity generation	Annual electricity generated, averaged over the asset lifetime	$kWh/year$
$X$	Peak-capacity	Peak power capacity of a newly built wind/solar PV asset	$kW$

$\gamma$	O&M cost over initial investment cost	Simplified notation for $\gamma = O/\mathcal{I}$	$\%/year$
$\mathcal{M}$	Module (solar PV) or turbine (wind) investment cost	Initial investment costs for the solar modules or wind turbines of a new asset, per peak capacity	$USD/kW_{peak}$
$\mathcal{A}$	Set of LCOE components	This set includes LCOE components $A \in \mathcal{A} = \{\mathcal{I}, \mathcal{C}, O, \mathcal{Y} + \gamma\}$	n.a.
$\mathcal{E}$	Average LCOE error	Multiplicative error term between the average national LCOE and the LCOE of the national averages	n.a.
$Z$	Cumulative renewable capacity	Total capacity (global or national) of solar or wind	$kW$

Table B.1: Overview of notation used throughout this article.

## B.2 Description of national data collection and cleaning

### B.2.1 Capacity and generation data

*Capacity* denotes the AC capacity of an electricity generation asset connected to the public or private grid (in kW). *Generation* denotes the electricity generated by this asset (in kWh). We collected annual data on the operational capacity, as well as the total electricity generated by this capacity. We did so by consolidating different datasets, particularly those of [BP \(2022\)](#), [IEA \(n.d.\)](#), and [IRENA \(2022\)](#). While these sources are mostly consistent with each other, their geographical coverage differs. Wherever these sources disagree, we have selected the data points most consistent with other data points and other reports from the literature.

### B.2.2 LCOE data

The *levelized cost of electricity (LCOE)* denotes the cost of a single asset to produce one unit of electricity, averaged over the asset’s lifetime. We consolidated a dataset of (capacity-) weighted averages of national LCOE. This data was obtained by manually

scraping the extended scientific and grey literature.

While this diverse set of sources increases the available data, it also increases measurement noise and inconsistencies. In particular, different sources may use different assumptions to clean their dataset. We correct these inconsistencies where possible, as described below.

### **LCOE data corrections**

Since we are interested in the LCOE at the year of installation, the LCOE is an ex-ante estimation over the asset lifetime. This means that LCOE is frequently calculated using synthetic and uniform assumptions about operating and maintenance costs, electricity production, and capital recovery. The LCOE data can be inconsistent across sources if these assumptions differ. We perform several corrections on the LCOE data. These corrections are largely modelled on IRENA's data ([IRENA, 2024](#)).

**Cost of Capital** The *cost of capital* or *weighted average cost of capital (WACC)* denotes the annual returns a project developer needs to pay debt and equity holders for financing the initial total investment costs. Before 2021, IRENA applied simplified cost of capital assumptions to its project-level data ([IRENA, 2023](#)). While the methodology has been significantly improved since then, previous years have not been updated. We, therefore, need to correct the cost of capital before 2021 to include country- and technology differentiation. After 2021, IRENA applied country- and technology-specific WACCs. These may still differ from the WACC we collected. To avoid artificially increasing the noise in our model based on these explicit assumptions, we harmonize the LCOE data to be consistent with our WACC data.

**Capacity Factor** The *capacity factor* denotes the yield factor between electric capacity and the average annual electricity generation of an asset. In other words, it determines how much electricity (kWh) is generated per capacity (kW) in each year, averaged over the asset lifetime. When calculating the LCOE of a newly commissioned renewable electricity

asset, IRENA relies largely on the capacity factor reported by the project developers. As discussed in Section B.3.7, this expected capacity factor can differ from the realized capacity factors used in our analysis. The reported capacity factor frequently exceeds what we observe in the market. This is particularly the case for wind electricity, where higher capacity factors from increased rotors and hub heights can have less impact on electricity generation than expected. We, therefore, correct the capacity factor for wind using the reported national average capacity factor. For solar, the reported capacity factors do not differ significantly from the ones calculated by our method and do not need to be corrected.

**Operating and Maintenance cost** The *operating and maintenance cost* (O&M) denotes the annual expense of producing electricity with a renewable asset, per operational capacity, averaged over the asset lifetime. While IRENA provides detailed information on O&M costs post 2010, they do not give details on costs prior to 2010.

To get a better understanding of the O&M costs used by IRENA, we can use Equation (1) in the main text to reverse-engineer their O&M assumptions:

$$O_{ij} = \mathcal{L}_{ij}\mathcal{Y}_{ij} - \mathcal{I}_{ij}\mathcal{C}_{ij}, \quad (\text{B.1})$$

Since the resulting O&M costs differ from our findings in Section B.3.8, we replace these costs with our assumptions.

An important caveat of such data corrections is that they limit our ability to test the forecasts made on corrected rather than raw data. If we use the same data to correct our LCOE as the one that we are making our forecasts with, we can not investigate the forecasting errors caused by this component of the LCOE. This particularly concerns the WACC and O&M cost. The Capacity Factor is not effected since it is based on capacity and generation data that is not used directly in our forecasting model. We, therefore, do not make stochastic forecasts for WACC and O&M costs but treat these as exogenous.

For the capacity factors, we can make stochastic forecasts based on actual generation and capacity data but can not investigate the potential correlations of noise with other components of the LCOE.

### **Different estimation methods for LCOE**

Different data providers may estimate the LCOE differently. For example, the LCOE can be estimated from Equation (1) in the main text or based on interviews with industry experts. Even within Equation (1), we observe differences, such as some sources applying a discount rate to the electricity generation to represent a preference for electricity today vs. in the future.

In our dataset, we do not distinguish between these methods and do not correct for any resulting differences. While this significantly increases the available data points, it also increases the associated noise. We have also chosen to use proxies for LCOE in some cases, namely the values of commercial Power-Purchase Agreements (PPAs), governmental (all-in) Feed-In Tariffs (FITs), and Contract for Difference (CfD) auction strike prices. Both of these represent the price of renewable electricity rather than the cost of electricity generation. However, we have found that the deviations between PPA/FIT and LCOE are usually smaller than the noise within different LCOE estimations (within a given country and year). That does not mean that these values are generally substitutable. For example, early FITs are frequently lower than the actual LCOE due to the political reality of these policy measures. We have, therefore, triangulated different estimation methods to exclude extreme values and inconsistent data points. This leaves us with 669 data points for solar PV and 1002 data points for wind. While this is still a relatively small number of data for our purpose of statistical forecasting and testing, it is significantly larger than what was previously used in the literature. For example, [Elshurafa et al. \(2018\)](#) analyze 347 data points for national solar BOS costs (although they differentiate commercial and residential solar, almost doubling their dataset), [Kothari et al. \(2023\)](#) use 112 data points for national BOS costs, [Riva et al. \(2018\)](#) analyze 52 data points for wind O&M costs.

### B.2.3 Total investment costs

*Total investment costs* denote the total overnight costs of purchasing and installing a new renewable electricity asset. It includes the costs for equipment, labour, permitting, land, grid connection, and other soft costs. It typically does not include the cost of integrating that asset into the public grid by providing additional transmission and distribution lines, as well as system flexibility to balance out variable renewable generation. As with LCOE, we consider the (capacity-) weighted average total investment cost per capacity. The process of collecting data on total investment costs is similar to that of LCOE. The only difference is that we are not using proxies for this data collection and do not make any corrections to the raw data. Unlike LCOE data, investment costs are more commonly reported by project developers and rely on fewer assumptions.

### B.2.4 Module and Turbine costs

*Module costs* denote the total costs of purchasing a new solar PV module, *Turbine costs* the costs for a new wind turbine. It includes the costs of production, transport, import tariffs and supplier margin. We consider the capacity-weighted average cost in both cases. PV module costs are generally reported per DC capacity. We can use the inverter loading ratio to convert this to AC capacity to make them comparable to wind turbine costs.

The national cost data for PV modules is based on the annual *Trends in Photovoltaic Applications* Report by the IEA's Photovoltaic Power Systems Programme (PVPS) (Masson and Kaizuka, 2022). Since national module costs are not reported consistently, we combine data from the different annual reports going back to 2002. To obtain the longest possible global time series, we use globally aggregated data from Our World in Data (OWID) that combines different sources since 1975 (*Solar (Photovoltaic) Panel Prices*, n.d.).

Turbine data is based on IRENA (2021), who combine different turbine selling prices depending on the market. We calculate the global weighted average cost by considering

the share of Western (Europe and the US, largely dominated by Vestas turbines) and Chinese manufacturers over time. The manufacturing share is based on granular studies of the turbine market (Gosens and Lu, 2014; Wood Mackenzie, 2024).

### B.2.5 Balance of system cost

Unlike capacity, generation, LCOE and total investment costs, balance of system (BOS) costs are rarely observable directly, apart from some exceptions such as Elshurafa et al. (2018). This means that we need to estimate the BOS costs from total investment costs and module/turbine costs. Using the equation from the main text, we can write

$$\mathcal{I}_{jt} = \mathcal{M}_{jt-\tau} + \mathcal{J}_{jt}. \quad (\text{B.2})$$

The factor  $\tau \geq 0$  is a lag between the procurement of modules/turbines and the procurement of BOS costs. For solar modules, we do not need to pay special attention to this, with  $\tau \approx 0$  since project planning and development typically only requires a few months (Beattie, 2021; Cruce et al., 2022). For wind turbines, we do need to consider project development timelines. Turbine procurement is typically conducted 1-2 years before an asset's operation date (U.S. Department of Energy, 2022). We, therefore, need to take a procurement time difference of  $\tau = 1.5$  into account. We smooth the reported turbine costs over time and apply this lag to calculate  $\mathcal{J}$ . As we show below, our results are not sensitive to this assumption in  $\tau$ .

### B.2.6 Capacity factor

The capacity factor  $\mathcal{Y}_{ij}$  gives the relationship between a renewable asset  $i$ 's electricity output and the installed capacity. We can define it as

$$\mathcal{Y}_{ijt} = \frac{Y_{ijt}}{X_{ijt}}, \quad (\text{B.3})$$

where  $Y_{ijt}$ (kWh/year) is the average annual electricity production of the asset with nameplate capacity  $X_{ijt}$ (kW) that started operation in year  $t$ .

Since the capacity factor is an ex-post average over an asset's lifetime, it needs to be estimated ex-ante at the beginning of a renewable asset's life. This means that there are different ways of estimating an asset's capacity factor. For example, a wind farm's capacity factor can be modeled based on the expected wind speeds in the farm's location or based on historical capacity factors for similar wind farms.

In our case, we can estimate the capacity factor either from the LCOE data we collected or from the capacity and generation data in our dataset. For wind, we find that the method based on the LCOE is not reliable. We thus use the electric capacity and generation to estimate the capacity factor. For solar, we find that both methods are comparable, with slightly lower errors from estimating the capacity factor based on the LCOE. The details of this are explained in Section [B.3.7](#).

## B.2.7 Operating and maintenance costs

Operating & Maintenance costs  $O_{ijt}$  are the costs accrued over the lifetime of a renewable project after the initial investment. They include maintenance and repair costs, operation management, commercial management, and site management ([Steffen et al., 2020](#)). Compared to fossil-fuel operations, operating costs are a small contribution to the LCOE for renewables. As a result, Operating & Maintenance costs have received significantly less attention than other cost components, and the available empirical data is limited ([Steffen et al., 2020](#); [Wiser et al., 2019, 2020](#); [IRENA, 2024](#); [Riva et al., 2018](#)).

Estimations for Operating & Maintenance costs between these sources differ substantially. For example, IEA reports estimate O&M costs at around 3.0 - 3.5% of overnight investment costs in Germany in 2012 ([Riva et al., 2018](#); [Comission de regulation de L'energie, 2014](#)). while [Raupach-Sumiya et al. \(2015\)](#) calculate almost 11%. Here, we consolidate these sources and make our own estimates of O&M costs relative to the total investment costs.

## B.2.8 Capital recovery factor

The Capital Recovery Factor (CRF) is determined by the after-tax weighted average cost of capital  $\mathcal{D}_{ijt}$  (or WACC) and the asset's lifetime  $T_i$ . The WACC represents the share of the initial investment  $\mathcal{I}_{ijt}$  that covers the expected annual returns of equity- and debt-investors. We calculate the CRF from the WACC as [Egli et al. \(2018\)](#),

$$\mathcal{C}_{ijt} = \frac{1}{\sum_{t=1}^{T_i} \frac{1}{(1+\mathcal{D}_{ijt})^t}} = \frac{(1 + \mathcal{D}_{ijt})^{T_i} \mathcal{D}_{ijt}}{(1 + \mathcal{D}_{ijt})^{T_i} - 1}. \quad (\text{B.4})$$

Furthermore, we can write the country-average WACC  $\mathcal{D}_{jt}$  for a technology  $k$  and country  $j$  as a sum of the country-specific risk-free rate  $r_{jt}$  and technology-specific risk premium  $p_k$  ([Polzin et al., 2021](#)):

$$\mathcal{D}_{jt}^k = r_{jt} + p_k. \quad (\text{B.5})$$

We can infer the risk-free rate from macroeconomic data. Specifically, we use long-term (10 and 7-year) government bond yields from [Bloomberg \(2022\)](#) and [Reuters \(2022\)](#). For countries without data, we use proxies from [IMF \(2022\)](#). In rare cases (for example, years with government defaults), we infer bond yields based on previous years and neighbouring countries.

Technology risk premia  $p_k$  may be subject to learning and could decline over time. The literature, however, diverges about the size of that effect. [Egli et al. \(2018\)](#) suggest that financing experience contributed around 1% of solar PV and 4% of onshore wind cost reductions in Germany between the 2000-2005 average and 2017. [Kempa et al. \(2021\)](#) attribute all cost-of-debt reductions to more stringent environmental policies and the maturity of the national banking sector instead of learning. While learning effects are likely, they appear to be small and not significant with respect to our stochastic forecasts. We therefore consider  $p_k$  as fixed over time.

[Steffen \(2020\)](#) provides one of the most comprehensive literature reviews and country- and technology-specific estimates of the after-tax WACC for renewables of 47 different

countries. Unfortunately, despite our simplifying assumption, the data is not available for all countries in our sample. We, therefore, need to infer  $p_k$  for other countries. We follow [Polzin et al. \(2021\)](#) that estimates WACC across Europe using linear interpolation between geographically close locations and a fixed ratio between solar and wind markups. This can further be applied to North America and East Asia. For Africa, we apply the same method based on data from [Thornton \(2018\)](#). For South America and the Middle East, where data is particularly limited, we use the theoretically derived data in [Ondraczek et al. \(2013\)](#).

There are two important assumptions underlying equations (B.4) and (B.5). Firstly, we assume that an asset’s lifetime is equal to the financing period. This assumption is adequate from an investor’s perspective but has some limitations. In particular, a social planner may want to differentiate the project lifetime and financing lifetime if the two differ substantially. In such applications, we can re-scale the CRF to the appropriate WACC and asset lifetime. Secondly, we assume that renewable financing is done through individual projects. This means that the WACC is independent of corporate finances unrelated to a specific renewable project. [Steffen \(2018\)](#) shows that this is appropriate for current and historical projects.

### B.3 LCOE model development

The model development consists of two key steps:

1. Decompose the asset-level LCOE  $\mathcal{L}_{ijt}$  into its components  $A_{ijt} \in \{\mathcal{I}_{ijt}, \mathcal{C}_{ijt}, \mathcal{O}_{ijt}, \mathcal{Y}_{ijt}\}$ .
2. For each component  $A$ , develop a model that outputs forecasts for  $A_{jt}$ , the capacity-weighted average of  $A_{ijt}$  over assets  $i$  in country  $j$  and year  $t$ .

We describe each step in detail, including the model validation performed to confirm the reliability of our model outcomes and reject model alternatives.

We first describe the LCOE decomposition (Step 1). The subsequent sections detail the component models (Step 2). Sections [B.3.2](#) and [B.3.3](#) describe the PV module and wind

turbine model development, respectively. Sections [B.3.4](#) and [B.3.5](#) are significantly more involved in developing our models for solar and wind BOS costs. In particular, we explore the role of local learning in BOS costs and how to use it in forward-looking forecasts. We aggregate module, turbine, and BOS costs to total investment costs in Section [B.3.6](#). Section [B.3.8](#) provides additional details on the operating and maintenance cost model.

### B.3.1 LCOE decomposition and aggregation

**Summary:** *We decompose the non-linear LCOE into its components using a linear approximation with a new stochastic error term. We use a normal distribution to model this term. For each component, we make use a statistically validated forecasts. Using an out-of-sample test, we further show that the aggregate LCOE model does not overfit the empirical data.*

Our starting point is the asset-level LCOE  $\mathcal{L}_{ijt}$ . Re-stating the LCOE decomposition from the main text, it is given as

$$\mathcal{L}_{ijt} \equiv \frac{\mathcal{I}_{ijt} * \mathcal{C}_{ijt} + \mathcal{O}_{ijt}}{\mathcal{Y}_{ijt}} \quad (\text{USD} / \text{kWh}), \quad (\text{B.6})$$

for asset  $i$  in country  $j$  and year  $t$ . We are ultimately interested in the capacity-weighted average LCOE  $\mathcal{L}_j$  for each country  $j$ , specific to the technology and year. Each component model forecasts future values of the capacity-weighted average component  $(\mathcal{I}_j, \mathcal{C}_j, \mathcal{O}_j, \mathcal{Y}_j)$ . Unfortunately, calculating  $\mathcal{L}_j$  from the capacity-weighted average components is not straightforward due to the non-linear nature of Equation (B.6). The national weighted average of the LCOE may differ from the LCOE of the weighted average component.

We, therefore, perform a log-linear approximation of equation (B.6) that we can use to make approximate forecasts for  $\mathcal{L}_j$ . This generates an additional error term  $\log \mathcal{E}$ , which we quantify empirically. Afterwards, we compare the disaggregate model to the global Wright’s law model of [Lafond et al. \(2018\)](#) to confirm that we have not over-fit our data.

## Linear LCOE approximation

In this section, we assess the relationship between the weighted average LCOE  $\mathcal{L}_j$  and the weighted average components  $\{\mathcal{I}_j, \mathcal{C}_j, O_j, \mathcal{Y}_j\}$ , particularly:

$$\mathcal{L}_j \stackrel{?}{=} \mathcal{L}(\mathcal{I}_j, \mathcal{C}_j, O_j, \mathcal{Y}_j). \quad (\text{B.7})$$

We first consider the LCOE of a single asset  $i$ , dropping the subscript  $t$  for convenience. As discussed in the main text and section B.3.8, we model the operating and maintenance costs as a fixed share  $\gamma$  of the total investment costs, i.e.,  $O_{ij} = \gamma * \mathcal{I}_{ij}$ . This means we can re-write the asset-level LCOE (B.6) as a linear equation:

$$\log \mathcal{L}_{ij} = \log \mathcal{I}_{ij} + \log (\mathcal{C}_{ij} + \gamma) - \log \mathcal{Y}_{ij} = \sum_{A \in \mathcal{A}} \log A, \quad (\text{B.8})$$

for components  $A \in \{\mathcal{I}, \mathcal{Y}^{-1}, O, \mathcal{C} + \gamma\} =: \mathcal{A}$ . Deriving the relationship between weighted average LCOE and weighted average components thus reduces to assessing the difference between the weighted average of a variable and the weighted average of its logarithm.

Based on Jensen's inequality, we know that the log of the average is greater or equal to the average of the log. Since this applies to both sides of the equation (the LCOE  $\mathcal{L}_{ij}$  and components  $A_{ij} \in \mathcal{A}$ ) the resulting biases cancel each other out to some degree but not necessarily completely. We can describe this relationship explicitly using a first-order Taylor expansion around the average LCOE and components.

Without going into the details of the Taylor expansion, we can write the average of the logarithm of a random variable  $X$  as the logarithm of the average,  $\bar{X}$ , such that

$$\overline{\log X} \approx \log \bar{X} + \varepsilon_X, \quad (\text{B.9})$$

where the first-order error term  $\varepsilon_X$  depends on the variance in the random variable  $X$ . In particular, if  $X$  has a low variance,  $\varepsilon_X$  is small, but if the variance is high,  $\varepsilon_X$  is large.

Using this identity, we can write Equation (B.8) as

$$\log \mathcal{L}_j \approx \log \mathcal{I}_j + \log(\mathcal{C}_j + \gamma) - \log \mathcal{Y}_j + \varepsilon_{\mathcal{I}} + \varepsilon_{\mathcal{C}} - \varepsilon_{\mathcal{Y}} - \varepsilon_{\mathcal{L}} \quad (\text{B.10})$$

$$\approx \log \mathcal{I}_j + \log(\mathcal{C}_j + \gamma) - \log \mathcal{Y}_j + \log \mathcal{E}_j, \quad (\text{B.11})$$

where  $\log \mathcal{E}_j$  is an aggregate first-order error term from the Taylor expansion.

Unfortunately, we cannot decompose  $\log \mathcal{E}_j$  further since this requires asset-level data on all LCOE components. In future research, asset-level cost data could be used to better capture the variation around the component average. This is particularly helpful in assessing how policies targeted at different sub-markets of the renewables industry could lower the average national LCOE.

### Aggregate error model

We can, however, study  $\log \mathcal{E}_j$  empirically as an aggregate error term. In other words, we investigate the difference between  $\log \mathcal{L}_j$  and  $\log \mathcal{I}_j \mathcal{Y}_j^{-1} (\mathcal{C}_j + \gamma)$  from Equation (B.11). We first perform a regression analysis on Equation (B.11). This tells us how much of the LCOE variance is explained by each component and  $\log \mathcal{E}_j$ . We then investigate the trend and correlation of  $\log \mathcal{E}_j$ . We do not find significant evidence of either, implying that we can reasonably model  $\log \mathcal{E}_j$  as i.i.d. random variables. Based on our dataset alone,  $\log \mathcal{E}_j$  follows a fat-tailed Cauchy distribution. However, when we consider a more granular dataset for the US (Barbose et al., 2024), we find that most extreme values in  $\log \mathcal{E}_j$  are caused by underlying data issues. We, therefore, exclude extreme values and use a normal distribution to model  $\log \mathcal{E}_j$ .

**Regression analysis** We consider the model

$$\log \mathcal{L}_j = \beta_{\mathcal{I}} \log \mathcal{I}_j + \beta_{\mathcal{C}} \log(\mathcal{C}_j + \gamma) - \beta_{\mathcal{Y}} \log \mathcal{Y}_j + \epsilon_j. \quad (\text{B.12})$$

	Solar					Wind				
$\log \mathcal{I}$	-0.24			0.96	1.0	-0.23			1.0	1.0
$\log \mathcal{Y}$		-0.28		-1.33	-1.0		-0.24		-1.26	-1.0
$\log(\mathcal{C}_j + \gamma)$			0.83		1.1			0.93		0.99
$R^2$	0.69	0.79	0.81	0.97	0.99	0.90	0.93	0.95	0.97	0.99

Table B.2: Regression results for LCOE against its components, according to Equation (B.12). We omit the indication of statistical significance since all regressors are highly significant above 99%.

Table B.2 shows the results of the regression. Equation (B.12) explains 99% of the variance in the data, leaving  $\log \mathcal{E}_j$  explaining the remaining 1%. It's interesting to see that if we include only the investment cost  $\mathcal{I}$  and capacity factor  $\mathcal{Y}$ , we can already explain 97% of the variance. In this case, the capacity factor represents a country-fixed effect, while the trend in the LCOE is closely aligned with the trend in the investment cost. Surprisingly, the investment cost alone explains a significant share of the variance but with the opposite sign. This is likely caused by the significant cross-sectional variance and the fact that the statistical significance of the model parameters is most likely overstated since it does not account for the correlation in our data.

Figure B.1 shows  $\log \mathcal{E}_j$  for individual countries  $j$  over time. The size of  $\log \mathcal{E}_j$  is comparable for solar and wind. By definition, both distributions are centred around zero since we use a fixed effect model to estimate the capacity factor. The fact that Turkey and Argentina drive the standard deviation increases in 2021 and 2022. In both cases, very large interest rates drive up the CRF, which is not reflected in the LCOE data.

**Trend analysis** For the trend in  $\log \mathcal{E}_j$  we maximize the available data by calculating the first (year-on-year) differences in national data. Based on a two-sided t-test, we do not find a significant positive or negative trend. For the trend in the cross-sectional variance of  $\log \mathcal{E}_j$ , we also calculate its first (year-on-year) difference. Similarly, a two-sided t-test finds no significant positive or negative trend.

**Autocorrelation analysis** We calculate the autocorrelation function (ACF) for each country's  $\log \mathcal{E}_j$  time series and look at the average ACF across countries. The ACF of a

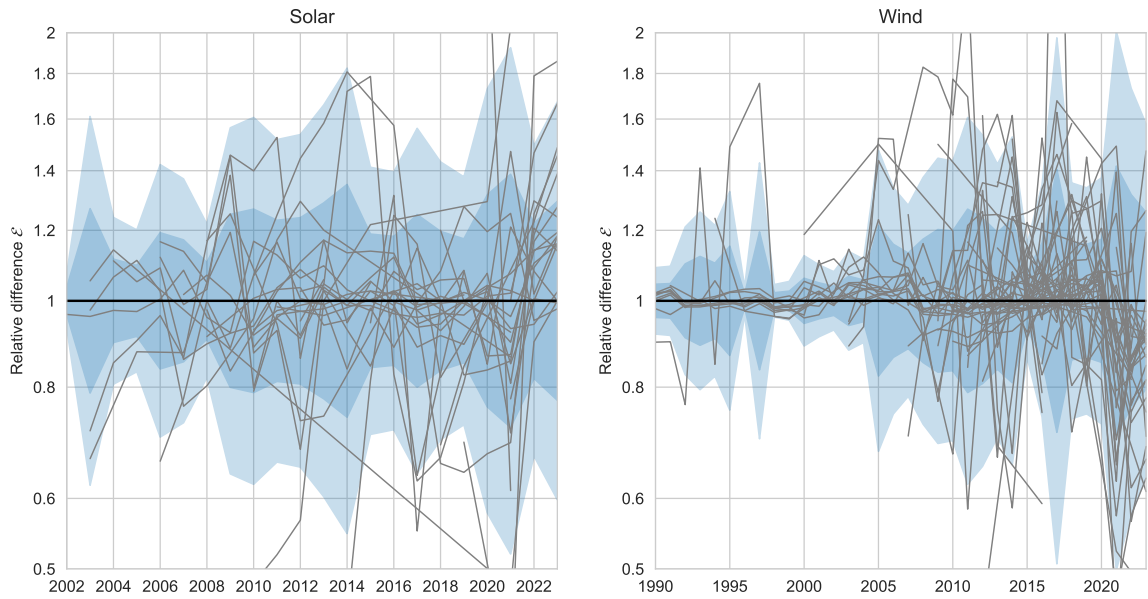


Figure B.1: **Natural Logarithm of the relative difference between the average LCOE  $\mathbb{E}_\omega \mathcal{L}$  and the LCOE of the averages  $\hat{\mathcal{I}}\hat{\mathcal{Y}}(\hat{\mathcal{C}} + \hat{\gamma})$  over time.** Grey lines show  $\mathcal{E}$  for a specific country over time. Blue shading indicates one / two standard deviations of the entire  $\mathcal{E}$  sample within one year. We observe that the distribution of  $\log \mathcal{E}$  is stationary over time, with no significant drift or changes in the standard deviation over time.

time series  $(X_t)_{t \in \mathbb{N}}$  and lag  $\tau$  is defined as the serial correlation between  $(X_t)_{t \in \mathbb{N}}$  and the lagged series  $(X_{t-\tau})_{t \in \mathbb{N}}$ . While the ACF for individual countries will likely be very noisy, the average thereof should be significantly less so and give us a meaningful indication of the correlation. The resulting ACF is shown in Figure B.2, with solar on the left and wind on the right panel. We can see that the ACF takes very low values below 0.25 for all tested lags. We conclude there is no significant correlation between individual years of  $\log \mathcal{E}_j$ , and we can sample them independently for our forecast.

**Distribution analysis** The previous findings imply that to forecast future values of  $\log \mathcal{E}_j$  we can reasonably approximate them by i.i.d. samples around zero. To define how we want to sample  $\log \mathcal{E}_j$ , we first consider the in-sample distribution based on our dataset.

The in-sample distribution is shown in Figure B.3, with solar on the left and wind on the right panel. We see that both distributions are symmetric around zero. Both distributions show extremely positive and negative values. This indicated that  $\log \mathcal{E}_j$  follows a distribution with two heavy tails. A Cauchy-Lorenz distribution with heavy

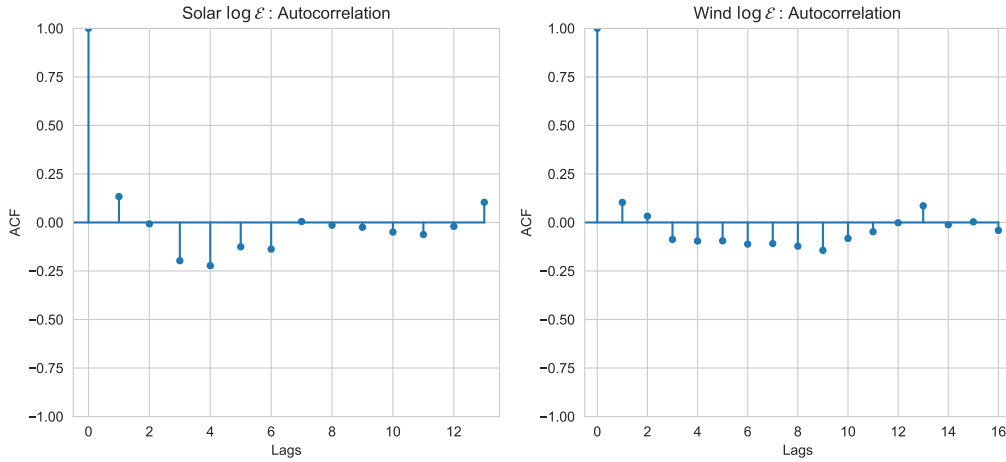


Figure B.2: **Average national autocorrelation function**  $\log \mathcal{E}$ . We take the average of the ACF of  $\log \mathcal{E}$  for individual countries for solar PV (**left**) and wind (**right**). Neither technology shows significant autocorrelation for the first lags.

tails fits this distribution well, as indicated by the orange line in Figure B.3. Similarly, a Cauchy distribution based on a maximum likelihood estimator can be used to forecast  $\log \mathcal{E}_j$ , with appropriate out-of-sample accuracy.

Nevertheless, the large errors and short-term effects raise the question of whether extreme values of  $\log \mathcal{E}_j$  are actually caused by economic effects or rather issues with the underlying datasets. While asset-level component data is, unfortunately, not widely available, we can study the log-difference effects for the investment cost in the US using an asset-level dataset from Barbose et al. (2024).

Figure B.4 shows the average and log-average total investment costs in the US over time. In most years, average and log-average are close to each other. However, in 2014, 2016, 2017, and 2022, the two values differ substantially. As described in the main text, the reason for this is individual renewable assets with large capacity and particularly low costs.

If we were to include these errors in our forecasts, the predicted errors would be heavily inflated from the fat tails of the Cauchy distribution. We, therefore, exclude  $\log \mathcal{E}_{jt}$  values whose absolute value exceeds 30%. The remaining errors are forecast with a normal distribution, based on a standard deviation estimate of roughly 10% for both solar and

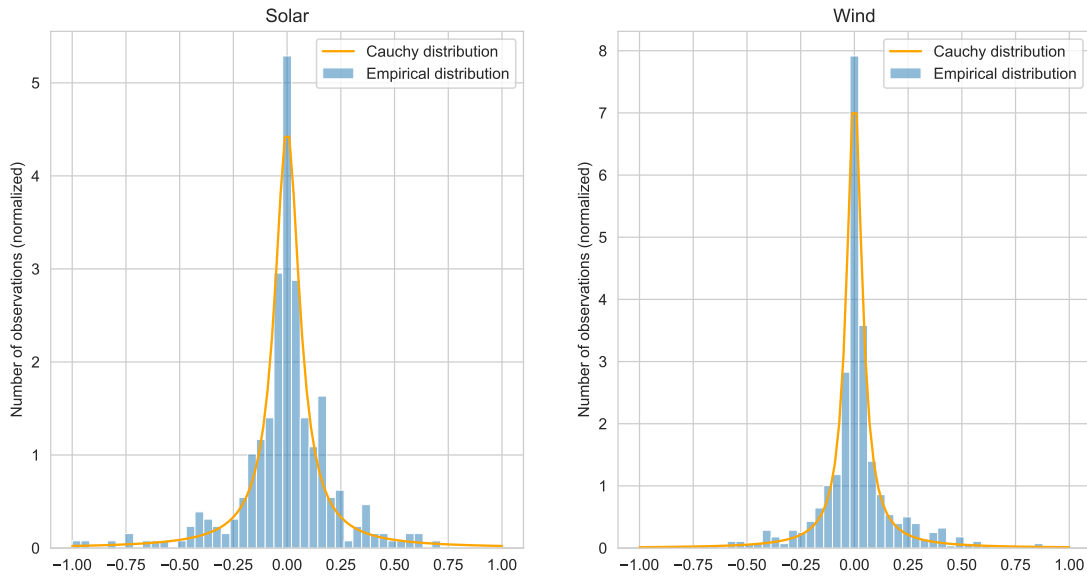


Figure B.3: **Empirical distribution of  $\log \mathcal{E}$** . We calculate  $\log \mathcal{E}$  for solar (**left**) and wind (**right**) and look at the resulting distribution. In both cases, we observe a large number of very positive or negative values of  $\log \mathcal{E}$ , indicating a heavy-tailed distribution.

wind.

### Forecast accuracy of the LCOE model

Conceptually, forecasting the LCOE from its components should work well since we include more information about the underlying dynamics. In practice, overfitting and parameter estimation errors in each component model are a real concern, so it may be better to forecast the aggregate directly (Grunfeld and Griliches, 1960). The choice between the disaggregate and a simpler top-down model depends on the sample size used to estimate parameters, the noise levels in the model, and the model dynamics (Lütkepohl, 1984). We answer this question based on empirical and theoretical considerations.

Empirically, we perform an out-of-sample cross-check between the aggregate model and simpler alternatives. The results of this are shown in Section B.3.9. The results are not statistically significant in either direction, indicating that we are not grossly overfitting our model.

We omit the theoretical consideration here since it is equivalent to the later discussion on total investment costs in section B.3.6. Since the investment costs are the only component




Figure B.4: **Average and log-average of the total investment cost  $\mathcal{I}$  in the US.** The black line shows the average  $\mathcal{I}_{jt}$  in the US over time, with the gray area indicating the 95th quantile of asset-level observation. The red dots show the exponential of the average of  $\log \mathcal{I}_{ijt}$  over time. For most years, it falls close to  $\mathcal{I}_{jt}$ , within the 95th quantile. There are, however, multiple outliers with an order of magnitude lower log-average costs.

subject to learning, with all other components being stationary or exogenous, we expect the disaggregate model to outperform the aggregate given sufficient data ( $\sim 25$  years) to estimate model parameters. Moreover, the disaggregate model provides much more information about the source of cost decline, making it the preferred choice in many applications.

### B.3.2 Module cost forecasts

**Summary:** *We model global average solar photovoltaic module costs and do not differentiate national costs. We use the stochastic Wright’s law model of Lafond et al. (2018) to forecast future global average module costs, conditional on global solar deployment.*

Most countries have had very similar module costs in the past (PVPS, 2015). This is expected with such a highly consolidated market; China has around 80% global market share in module manufacturing (Masson and Kaizuka, 2022). One exception to this is Japan. Protectionism in the Japanese market resulted in the decoupling of local prices from the global average after 2010 (Raupach-Sumiya et al., 2015). In line with most energy technology literature, we do not treat this decoupling explicitly but consider a global

module cost. We thus drop the country-subscript  $j$  from our notation for the remainder of this section.

For global average manufacturing costs, we can use an existing probabilistic forecast model. Lafond et al. (2018) show that one can model Wright’s law as a geometric random walk with nonlinear drift and autocorrelated residuals. This holds for a large number of technologies including solar modules. We therefore apply this model to project future solar module costs:

$$\text{Module model: } \log \mathcal{M}(t) = \log \mathcal{M}(t-1) - \hat{\omega} \log \frac{Z(t)}{Z(t-1)} + u_t + \rho u_{t-1}, \quad (\text{B.13})$$

where  $Z(t)$  is the global solar capacity until time  $t$ ,  $\omega$  is the estimated Wright’s law exponent with  $\hat{\omega} \sim \mathcal{N}(\omega, \hat{\sigma}_\omega^2)$ ,  $u_t \sim \mathcal{N}(0, \hat{\sigma}_u^2)$  is i.i.d. noise, and  $\rho = 0.19$  is the autocorrelation parameter. The parameters  $\hat{\omega}, \hat{\sigma}_\omega^2, \hat{\sigma}_u^2$  are estimated using a linear regression on the first difference of the module costs.

Since we do not alter the model used by Lafond et al., we do not need to perform a separate validation. The model in Equation (B.13) has been extensively studied on over 50 technologies, as well as applied to broader energy system scenarios (Way et al., 2022).

### B.3.3 Turbine cost forecasts

**Summary:** *We observe different costs for Chinese- and Western-manufactured wind turbines. However, due to the small sample size, we can not tell with statistical significance if a model that distinguishes manufacturing-location-specific learning is more accurate than a global model. We show that the stochastic Wright’s law model of Lafond et al. (2018) based on global learning provides appropriate uncertainty estimates.*

Unlike solar modules, wind turbines have shown substantially different costs in different locations. As shown in Figure B.5, Western-manufactured (Europe and the US) and Chinese-manufactured turbines used to have similar costs in the early 2000s. However,

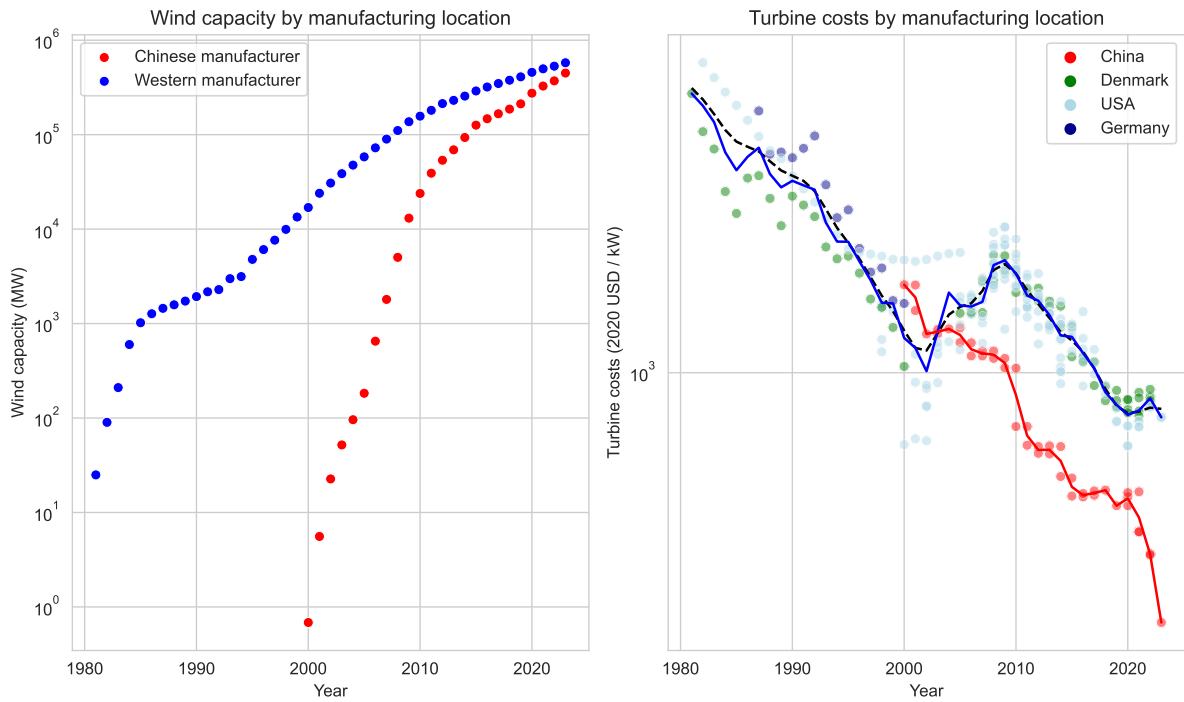


Figure B.5: **Wind turbine costs vs cumulative deployment and year.** The left panel shows the increase in wind turbine capacity manufactured by Chinese and Western companies between 1982 and 2021. The right panel shows the decrease in the respective turbine costs over the same period. The red and blue lines show a (weighted) average cost for Chinese/Western turbines. The black dotted line is a smoothed interpolation of the western turbine cost used to calculate wind BOS costs in subsequent sections.

since around 2005, Chinese turbines have shown significantly lower costs than their Western alternatives. While Western turbines have been around for longer, Chinese turbines have also grown more quickly. At the moment, they are mainly sold in China and, to some extent, India and Brazil, but foreign market share is continuously increasing, and Chinese turbines have recently overtaken Western alternatives in terms of global market share (IHSM, 2023; Gosens and Lu, 2014; Hayashi et al., 2018; Wood Mackenzie, 2024).

This raises the question of whether the lower costs of Chinese turbines are due to local experience effects and if we should differentiate between manufacturing locations in our forecasts. Unfortunately, with only two short time series, we do not have sufficient data to assess with statistical significance if a differentiated model is more appropriate. Nevertheless, we performed a cross-check in Section B.3.9 to compare model alternatives.

Nevertheless, we must select one of these four alternatives to make forecasts. Here,

we choose a global experience curve model in line with [Lafond et al. \(2018\)](#). In some applications, such as investigations of trade policies and disjoint markets, it may be more insightful to use a model with regional differentiation. Researchers will have to be careful to investigate the impact of their model choice on their results since we cannot clearly state a preference for one model over the other.

Since, unlike in [Lafond et al. \(2018\)](#), we use a global model to forecast regional turbine costs, we need to validate our model results. To do so, we need a full specification of the residuals associated with our forecasts. This way, we can identify whether our forecast residuals are consistent with the model.

We can use the same model as Equation (B.13) for each market  $j$ :

$$\text{Turbine model: } \log \mathcal{M}_j(t) = \log \mathcal{M}_j(t-1) - \omega \log \frac{Z(t)}{Z(t-1)} + \eta_t^j + \rho \eta_{t-1}^j, \quad (\text{B.14})$$

where  $\eta_t^j \sim \mathcal{N}(0, \hat{\sigma}_\eta^2)$  and  $\omega - \hat{\omega} \sim t_m * \hat{\sigma}_\omega$ . Here  $\hat{\sigma}_\eta$  denotes estimated variances of the noise,  $\hat{\omega}$  denote the estimated Wright's law exponent,  $\hat{\sigma}_\omega$  is the estimated standard errors of the parameters, and  $\omega$  is the true Wright's law exponent.  $t_m$  is a Student t-distribution with  $m$  degrees of freedom, equivalent to the number of data points used to estimate the model. We are also including the previous autocorrelation of our noise with  $\rho = 0.19$  ([Lafond et al., 2018](#)).

We estimate the model parameters  $\hat{\omega}$  and  $\hat{\sigma}_\eta$  with the same method as [Lafond et al. \(2018\)](#) on the global average turbine costs  $\mathcal{M}(t)$ . Since  $\mathcal{M}(t)$  is the weighted average between the two  $\mathcal{M}_j(t)$  time series, we need to correct for this effect when estimating the variance  $\hat{\sigma}_\eta$ . Since we have two time series following the same random process, we can do so with an additional factor of 2. So,  $\hat{\sigma}_\eta = 2 \cdot \tilde{\sigma}_\eta$ , where  $\hat{\sigma}_\eta$  is the estimated variance of the two  $\mathcal{M}_j$  processes, and  $\tilde{\sigma}_\eta$  is the estimated variance of their average  $\mathcal{M}$ .

Note that neither the Wright's law exponent  $\omega$  nor the exogenous shock scale  $\sigma_\eta^2$  are specific to the manufacturing market. Both markets share universal parameters since they follow the same learning dynamics. Secondly, [Smart and Farmer \(2024\)](#) show that in this

case of little data and similar model parameters, applying the same parameters to both technologies is preferred over technology-specific parameters.

We validate this model using backtesting. We consider all 5-year moving average windows between times  $t$  and  $t - 4$  to estimate our model parameters  $\Theta_t := (\hat{\omega}, \hat{\sigma}_\eta, \hat{\sigma}_\omega)$ . Based on these parameters, we make stochastic forecasts for the future of our training data  $t + \tau$ , using Monte Carlo sampling. This gives us the median point forecast and associated theoretical uncertainty distribution for each forecast horizon  $\tau$ . We can now compare the theoretical uncertainty distribution to the empirically observed residuals for each value of  $t$ ,  $\tau$ , and market  $j$ .

To compare these probabilistic forecasts to the true data, we compute the cumulative *probability integral transform* (PIT) of our forecasts (Gneiting and Katzfuss, 2014; Gneiting et al., 2007). The reason for using the cumulative PIT is to better examine statistical deviations thereof, as explained below. To calculate the PIT, consider a quantile  $p \in [0, 1]$  of our probabilistic forecast. If our quantiles are correct, we expect the share of empirical observations that lie below this quantile,  $q_\Theta(p)$ , to be equal to  $p$ , i.e.,  $q_\Theta(p) = p$ . Figure B.6 shows the observed quantiles  $q_\Theta(p)$  against  $p$ . The gray bars are the empirical observations. Theoretically, we expect them to follow the identity line indicated in red. Graphically, we see that the empirical observations are above the identity. This means that our probabilistic forecasts tend to be larger than the empirical observations instead of symmetrically scattered around them.

To assess if the empirical observations are significantly different from the theoretical model we apply a surrogate data method. We assume the same data structure as in our empirical observations and simulate 500 local geometric random walks conditional on future cumulative deployment, using the same Wright’s law parameters of our model. We then compute the cumulative PIT for each simulation run, i.e. estimate  $q_\Theta(p)$ . This provides confidence intervals for each value of  $p$ , shown as red areas in Figure B.6. Since for each surrogate dataset, the  $q_\Theta(p)$ -values are not independent from each other, we plot the cumulative PIT. This way, we can compute meaningful confidence intervals that help




Figure B.6: **Out-of-sample validation of the turbine cost forecast error.** We estimate model parameters on 5-year moving windows of our data. We make stochastic forecasts for the future of this window with a Monte Carlo approach. The empirical cumulative PIT  $q_{\Theta}(p)$  is shown in gray. The red line is the theoretically expected PIT. The red areas are the 50% and 95% confidence intervals under a surrogate random walk model. For almost all values of  $p$ , the empirical distribution falls within the 95% percentile, indicating that our model is a reasonable representation of reality. Between  $p = 0.5$  and  $p = 0.6$ , the empirical distribution is marginally outside of the 95% quantile but still well within the 96% quantile.

the graphical analysis. We see that our grey bars (the empirical data) fall just within the 95% confidence interval of the surrogate data, with one exception between the 50% and 60% quantile, which is within the 96% quantile. We conclude that there is no significant difference between our forecasts and an ideal model and that the model is appropriate.

While this shows that we can use the *Wright, Global (U)* model to make forecasts for turbine costs, it does not mean that the *Wright, Local (U)* model is inappropriate. In fact, when we repeat this exercise for the *Wright, Local (U)* model (not shown here), we also find that the observations fall within the respective 95% confidence interval.

### B.3.4 Solar BOS cost forecasts

**Summary:** *BOS costs display substantial national correlation. We study local and global learning models to forecast future national BOS costs. A global learning model is more appropriate than its local alternative due to the cross-sectional correlation. We validate a model that forecasts national deviation from the global average as a mean-reverting AR(1) process.*

To develop our solar BOS cost model, we first perform an in-sample analysis, followed by a similar validation exercise as for wind turbines. In addition, we illustrate how we have excluded other model alternatives.

### In-sample analysis

National solar BOS data display a significant temporal trend, cross-correlation between national BOS costs, and autocorrelation within each national BOS cost.

**Trend analysis** As evident from Figure 1 in the main text and confirmed by a panel regression analysis, there is a significant decline in the logarithm of the BOS costs with increasing cumulative deployment. We are thus looking for a learning model that includes this trend in our cost forecasts.

**Cross-sectional correlation analysis** We use a cross-sectional dependence test for unbalanced panel data, proposed by [Pesaran \(2021\)](#); [De Hoyos and Sarafidis \(2006\)](#) We estimate the statistic

$$CD = \sqrt{\frac{1}{N(N-1)}} \left( \sum_{j=1}^{N-1} \sum_{j'=j+1}^N \sqrt{T_{jj'}} \hat{\rho}_{jj'} \right), \quad (\text{B.15})$$

where  $N$  is the number of countries in our sample,  $j$  and  $j'$  are the country-pairs we iterate over,  $T_{jj'}$  is the number of years for which we have data in both countries and  $\hat{\rho}_{jj'}$  is the estimates pairwise Pearson correlation. Under the null hypothesis of no cross-sectional correlation,  $CD \sim N(0, 1)$ . We find that  $CD$  significantly differs from zero or that the national time series are significantly correlated.

If the cross-sectional correlation is large, the statistical significance of our analysis is compromised, even though this will not show up in many standard statistics packages. One way of assessing the limits of the statistical significance is to calculate the *Effective Number of Spatial Degrees of Freedom* (ESDOF), a measure originating in Meteorology

(Bretherton et al., 1999). The ESDOF shows that out of the 42 national BOS time series (59 for wind), only 9 (17 for wind) are statistically independent and can effectively be used for model validation.

To calculate the ESDOF for  $N$  time series, we normalize our data by taking the first differences and calculating the  $N \times N$  covariance matrix  $\mathbf{C}$ . Using a Singular Value Decomposition, we find the eigenvalues  $(\lambda_i)_{i \in [1, N]}$  of  $\mathbf{C}$ . The ESDOF  $N^*$  is then given by Bretherton et al. (1999)

$$N^* = \frac{(\sum \lambda_i)^2}{\sum \lambda_i^2}. \quad (\text{B.16})$$

Our model validation method has to reflect this correlation. For the remainder of this section, we deal with this challenge by subtracting the cross-sectional average BOS costs from the time series. Specifically, we calculate the national deviations as

$$\log \bar{\mathcal{J}}_j(t) = \log \mathcal{J}_j(t) - \log \mathcal{J}(t), \quad (\text{B.17})$$

where  $\log \mathcal{J}(t)$  is the unweighted global average across all  $\log \mathcal{J}_j(t)$ . The reason for using an unweighted average (compared to a capacity-weighted average) is that we do not want to prioritize accuracy for countries that deploy a lot of solar in our subsequent model validation. The resulting time-series  $\log \bar{\mathcal{J}}_j(t)$  are no longer cross-sectionally correlated, as confirmed by another cross-sectional dependence test.

**Autocorrelation analysis** We calculate the Autocorrelation function (ACF) and Partial Autocorrelation function (PACF) for each national time series. The ACF of a time series  $(x_t)_{t \in \mathbb{N}}$  for lag  $h$  is estimated as the Pearson correlation between all pairs  $x_t$  and  $x_{t-h}$ . The PACF of the time series is the correlation between  $x_t$  and  $x_{t-h}$ , correcting for any linear dependence on  $x_{t-1}, \dots, x_{t-h+1}$ . The ACF and PACF of the national time series are very noisy due to the limited number of data points in each country. To reduce the noise, we calculate a weighted median and confidence intervals of the (P)ACF. The weights are given by the length of each national time series since longer time series have a lower

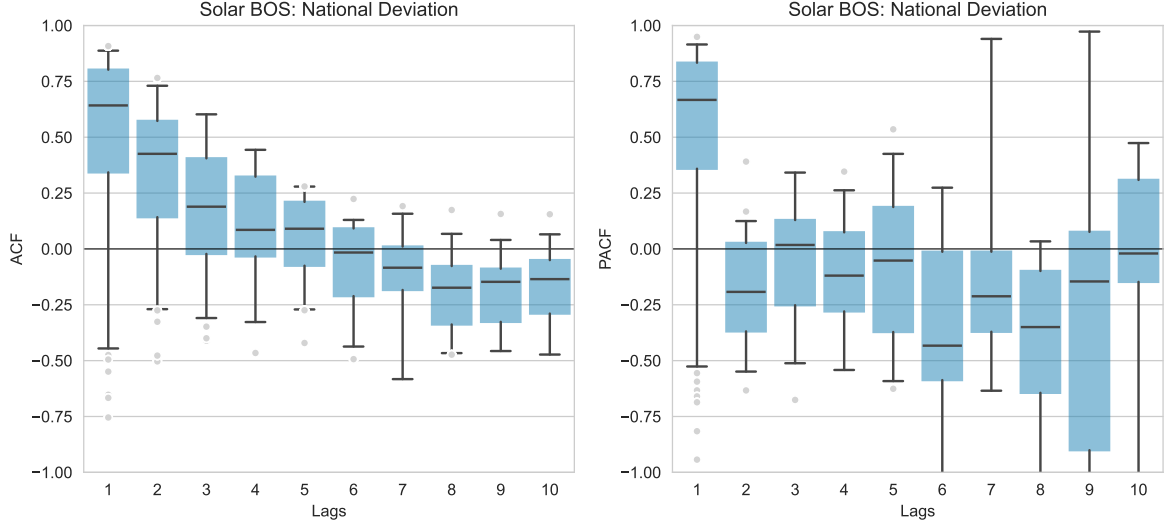


Figure B.7: **Cross-sectional weighted average ACF and PACF for the national solar BOS costs.** We calculate national ACF and PACF functions and take their cross-sectional median and confidence intervals. The weights are based on the length of each national time series. **(left panel)** The cross-sectional weighted average ACF plot. **(right panel)** The cross-sectional weighted average PACF plot. The line shows the (weighted) median ACF / PACF. The box shows the (weighted) interquartile range. The whiskers show the 90% confidence range. The dots show outliers exceeding this range. The ACF shows a marginally significant first lag. Similarly, the PACF shows a marginally significant first lag. These plots indicate a first-order autocorrelation.

measurement error for the (P)ACF. The resulting average functions are shown in Figure B.7. Both ACF and PACF show one significant lag. This indicates that our national time series are significantly autocorrelated.

### Solar BOS model validation

To validate the global forecasting model for solar BOS costs described in the main text, we require a full description of the errors associated with our forecasts.

$$\log \mathcal{J}(t) = \log \mathcal{J}(t-1) - \omega \log \frac{Z(t)}{Z(t-1)} + u_t + \rho u_{t-1}, \quad (\text{B.18})$$

$$\log \mathcal{J}_j(t) - \log \mathcal{J}(t) =: \log \bar{\mathcal{J}}_j(t) = \varphi \log \bar{\mathcal{J}}_j(t-1) + \eta_{jt}. \quad (\text{B.19})$$

We model the global average BOS costs  $\mathcal{J}(t)$  as a Wright's law model conditional on global solar capacity  $Z(t)$ , with i.i.d. normal shocks  $u_t \sim \mathcal{N}(0, \sigma_u^2)$ , Equation (B.18). The national deviations from the global average  $\log \bar{\mathcal{J}}_j$  are calculated by subtracting the global average




Figure B.8: **Cross-sectional variance of the logarithm of the solar BOS over time. (blue dots)** For each year, we compute the cross-sectional variance in the logarithm of national solar BOS costs,  $\log \bar{\mathcal{J}}_j(t)$ . The black line shows a linear regression of the resulting data. It shows a statistically significant decline.

from the national time series. The resulting time series is modelled with an AR(1) process, Equation (B.19). We estimate  $\varphi$  as the cross-sectional weighted average of estimated lag-1 correlation parameters, such that  $\varphi \approx 0.85$ . This is similar to the average lag-1 coefficient for wind, as discussed later. We estimate the scale of the local noise terms  $\eta_{jt} \sim \mathcal{N}(0, \sigma_\eta^2)$  based on the cross-sectional sample variance  $\hat{K}_\eta^2$ , similar to Farmer and Lafond (2016), where

$$\hat{K}_\eta^2 = \hat{\sigma}_\eta^2 / (1 - \varphi^2). \quad (\text{B.20})$$

As shown in Figure B.8, the cross-sectional variance  $\hat{K}_\eta^2$  is decreasing over time. Note that this variance is that of  $\log \bar{\mathcal{J}}_j$ , such that  $\hat{K}_\eta^2$  is relative and its decrease is not due to the decrease in BOS costs  $\mathcal{J}_j$ . One possible explanation of the variance decrease is our use of a universal solar module cost  $\mathcal{M}$  when calculating  $\mathcal{J}_j$ ; we do not account for national differences in module costs. This effect has been less pronounced in recent years due to the increasing consolidation of the module market. Nevertheless, we do not know if or how this trend will continue. Therefore, we model the national deviations conditional on the cross-sectional variation  $\hat{K}_\eta^2$  and assume that they remain constant in our forecasts.

We focus our model validation on the national deviations from the global average, i.e.,  $\log \bar{\mathcal{J}}_j$  in Equation (B.19) since the global model has been validated by Lafond et al. (2018). Equivalent to the turbine model described in Appendix B.3.3, we validate this

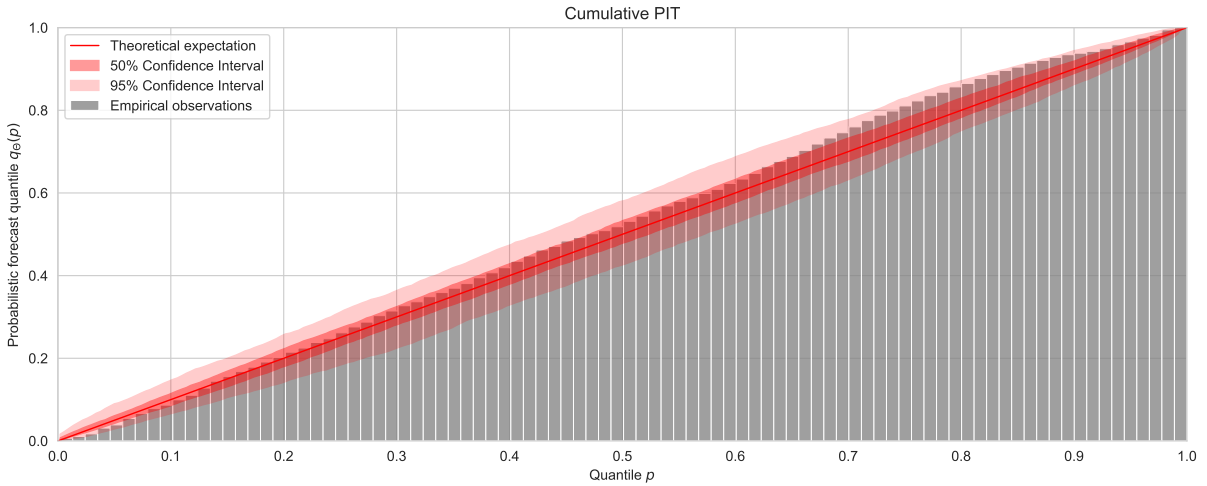


Figure B.9: **Out-of-sample validation of the solar BOS cost forecast  $\log \bar{\mathcal{J}}_j$ , relative to the true global average cost.** We consider only the national solar BOS cost relative to the global average BOS cost  $\log \bar{\mathcal{J}}_j$ . We make stochastic forecasts for the relative costs using the AR(1) model described in the main text, conditional on  $K_\epsilon^2(t)$ . The empirical cumulative PIT  $q(p)$  is shown in grey. The red line is the theoretically expected PIT. The red areas are the 50% and 95% confidence intervals under a simulated AR(1) model. We see that the forecast distribution is statistically identical to the AR(1) model.

model using backtesting. For each point in time, we make stochastic forecasts for future values of  $\log \bar{\mathcal{J}}_j$ , conditional on  $K_\epsilon^2(t)$ . We then calculate the cumulative PIT  $q(p)$  and compare it to what we would expect from a surrogate data simulation.

Figure B.9 shows the observed quantiles  $q(p)$  against  $p$ . We see that the empirical quantiles are close to the identity. While there are some deviations from the universal distribution, they are not statistically significant.

### Solar BOS model comparison

To illustrate our validation process, we further outline the empirical and theoretical analyses performed to exclude alternative forecasting models. We compare models that consider BOS costs a purely local component subject to local learning (Elshurafa et al., 2018), to models that consider BOS costs subject to global learning. Intuitively, we expect local BOS components to be independent from one another and subject to local learning. However, a global model may provide better forecasts for correlated time series, as we have here.

Table B.3 describes the four model alternatives we compare in this section. We consider

three Wright’s law models, where costs decline as a function of cumulative deployment, and one Moore’s law model, where costs decline over time. The *Wright, Global (U)* model is the chosen model described in equations (B.18)-(B.19). Global learning is the dominant effect, and national BOS costs depend strongly on what costs are observed in other countries. The other three models focus on local learning as the dominant effect, whereby future costs decline independent of other countries deployment. The difference between the *Wright, Local (U)* and *Wright, Local* model is that the former assumes the same learning rate across countries. While cumulative deployment may differ between countries, they all learn at the same rate.

**Wright, Local model** We reject the *Wright, Local* model based on one empirical and one theoretical argument. Empirically, we perform a backtesting exercise where the *Wright, Local* model performs worse than its alternatives. We then apply novel theoretical method (Smart and Farmer, 2024) to identify individual countries for which the *Wright, Local* model would outperform its alternatives. For around half the countries, the *Wright, Local (U)* model is preferred due to data limitations. For the remaining half, there is no statistically significant difference between the *Wright, Local* and *Wright, Local (U)* model. This means that we can just as well use the *Wright, Local (U)* model and reject the *Wright, Local* model.

We perform an out-of-sample test based on backtesting:

1. Estimate all model parameters on every 5-year window, i.e., between times  $t$  and  $t - 4$
2. Use the model parameters to forecast future BOS costs in times  $t + \tau \geq t$
3. Compare forecast residuals between models using an aggregate measure, the so-called *scoring rule* (Gneiting and Raftery, 2007); the model with the lowest scoring rule is considered the most accurate.

Here, we use the *Root Mean Squared Error* (RMSE) as our scoring rule. The RMSE of

Name	Description	Functional Form	Learning implications
Wright, Global (U)	We forecast the global weighted average $\mathcal{J}$ based on Wright's law, conditional on global cumulative deployment $Z(t)$ . We use the same forecast in all countries. The learning parameter $\omega$ is estimated using a linear regression on the global average. way22	$\log \mathcal{J}_j(t) = \log \mathcal{J}(t-1) - \omega \log \frac{Z(t)}{Z(t-1)}$	All countries learn with increasing global experience. There is significant learning spillover through shared experience, i.e., countries depend on each other's deployment.
Wright, Local (U)	We forecast each country independently, using Wright's law. We use the same universal Wright's law exponent $\omega$ for all countries based on a panel regression.	$\log \mathcal{J}_j(t) = \log \mathcal{J}_j(t-1) - \omega \log \frac{Z_j(t)}{Z_j(t-1)}$	All counties learn at the same rate, with increasing local experience. There is no spillover between countries and costs are independent of deployment in the other countries. Costs are likely to diverge if the cumulative deployment grows at different rates in different countries.
Moore, Local (U)	We forecast each country independently, using Moore's law. We use the same universal Moore's law exponent $\theta$ for both countries based on a panel regression.	$\log \mathcal{J}_j(t) = \log \mathcal{J}_j(t-1) - \theta$	All counties learn at the same rate, independent of their local experience. There is no spillover between countries and costs are independent of other countries. Costs are likely to diverge across countries over time.
Wright, Local	We forecast each country independently, using Wright's law. We estimate a country-specific Wright's law exponent $\omega_i$ for each country based on a regression model. elshurafa18	$\log \mathcal{J}_j(t) = \log \mathcal{J}_j(t-1) - \omega_j \log \frac{Z_j(t)}{Z_j(t-1)}$	All counties learn with increasing local experience but at different rates. There is no spillover between countries; costs are independent of deployment in the other countries. Costs are likely to diverge even if the cumulative deployment grows at the same rate across countries.

Table B.3: **Overview of compared forecast models for solar BOS costs.** We compare 4 different forecast models forecasting future solar costs in different countries.  $\theta$  and  $\omega$  are Moore's and Wright's law parameters, respectively, independent of the market.  $\omega_j$  is a Wright's law parameter specific to one market  $j$ . We choose the *Wright, Global (U)* model as most appropriate for our forecasts. We reject the other models on an empirical and a theoretical basis

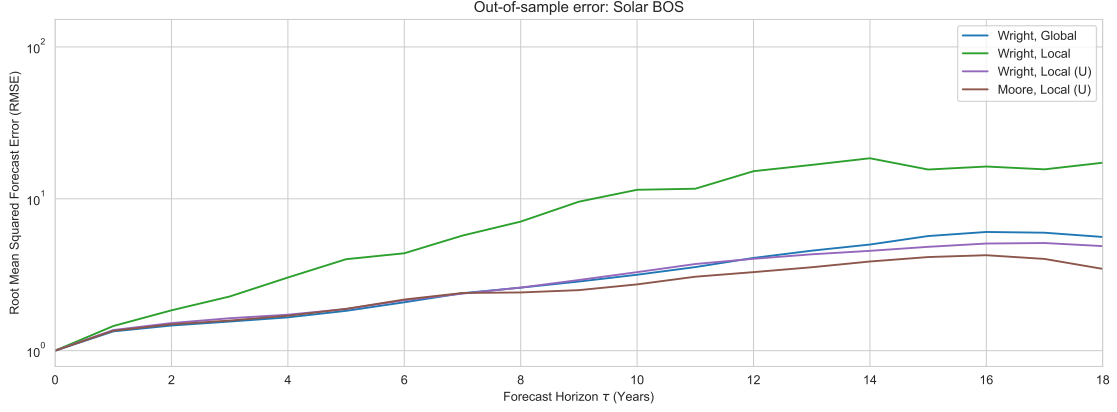


Figure B.10: **Out-of-sample residuals for different forecasting models of solar BOS.** BOS costs are forecast under different models described in Table B.3. We estimate the model parameters on 5-year moving average training data and compare the forecasts with the actual data. The x-axis shows the forecast horizon. The y-axis shows the RMSE of the logarithm of the BOS costs for our testing samples. The RMSE is shown as the relative error on the (linear) BOS cost. Except for the *Wright, Local* model, all models show similar forecast accuracy.

model  $l$  and forecast horizon  $\tau$  is defined as

$$RMSE(\tau) = \sqrt{\frac{1}{n} \sum_{t,j} \left( \log \hat{\mathcal{J}}_j^l(t + \tau) - \log \mathcal{J}_j^l(t + \tau) \right)^2}, \quad (\text{B.21})$$

where  $n$  is the number of times  $t$  and respective country  $j$  in our sample, and  $\hat{\mathcal{J}}_j^l(t + \tau)$  is the forecast of  $\mathcal{J}_j(t + \tau)$  for model  $l$ .

The results of this analysis are shown in Figure B.10. The *Wright, Local* model performs visibly worse than the other models. This difference is significant at 99% based on an F-test, Barlett-test, and Levene-test. However, none of these tests account for the correlation in our data. These results, therefore, need to be treated with care. Furthermore, the test results are sensitive to the size of the training data (5 years). For longer training data, we expect the difference in RMSE between the *Wright, Local* and *Wright, Local (U)* to approach zero or switch sign (Smart and Farmer, 2024; Pesaran et al., 2024). The reason for the lower accuracy of the *Wright, Local* model is that this model is prone to over-fitting—an issue that decreases with a larger training data set. Based on this empirical test, it is not clear if the *Wright, Local (U)* model still performs better than the *Wright, Local* model if we use all available data points to estimate the model parameters.

Since we cannot perform a backtesting exercise whilst also estimating our model parameters on all available data, we rely on the theoretical investigation by [Smart and Farmer \(2024\)](#) that compare geometric random walk forecast models with and without parameter heterogeneity. Under the assumption that the national learning parameters  $\omega_i$  follow a gamma distribution with parameters  $\alpha_\omega$  and  $\beta_\omega$ , each random walk is subject to i.i.d. noise shocks  $\epsilon_{it} \sim \mathcal{N}(0, \sigma_\epsilon^2)$ , and each parameter is estimated based on  $m$  national data points, the *Wright, Local (U)* model has lower forecast errors than the *Wright, Local* model iff  $m > \sigma_\epsilon^2 \beta_\epsilon^2 / \alpha_\epsilon$ . In our case, we find that  $\sigma_\epsilon \approx 0.42$ . This noise level is significantly higher than in the global model due to national fluctuations. For the distribution of  $\omega_i$ , we find that a gamma-distribution with  $\alpha_\omega \approx 27.0$  and  $\beta_\omega \approx 47.9$  describes observed in-sample national learning parameters well (excluding three data points with  $\omega_i$  exceeding 1.5). This means that conditional on *Wright, Local* being the true model for  $m \geq 15$ , the *Wright, Local* model is appropriate to forecast future national BOS costs.

Figure [B.11](#) shows the national parameters  $\hat{\omega}_i$  against the number of data points  $m$  available for each parameter estimation. The green shaded area indicates  $m \geq 15$ , i.e., parameters for which the *Wright, Local* model is preferred. We observe that for those parameters with  $m \geq 15$ , the universal parameter  $\hat{\omega}$  is always included in the 90% confidence interval of  $\hat{\omega}_i$ . This has important practical implications since it means that, for those cases where a national parameter is justified, it ceases to make a large difference in our forecasts. This means that in both cases,  $m < 15$  and  $m \geq 15$ , the *Wright, Local (U)* model is appropriate and statistically equivalent to the *Wright, Local* model when a switch could be made. Based on this, we exclude the *Wright, Local* model from further analysis.

**Wright, Local (U) & Moore, Local (U) model** There are two empirical arguments we use to reject the *Wright, Local (U)* model. Since these are equivalent for the *Moore, Local (U)* model, we only provide details on the *Wright, Local (U)* model. Firstly, we consider the full stochastic forecast that this model implies. Based on a backtesting

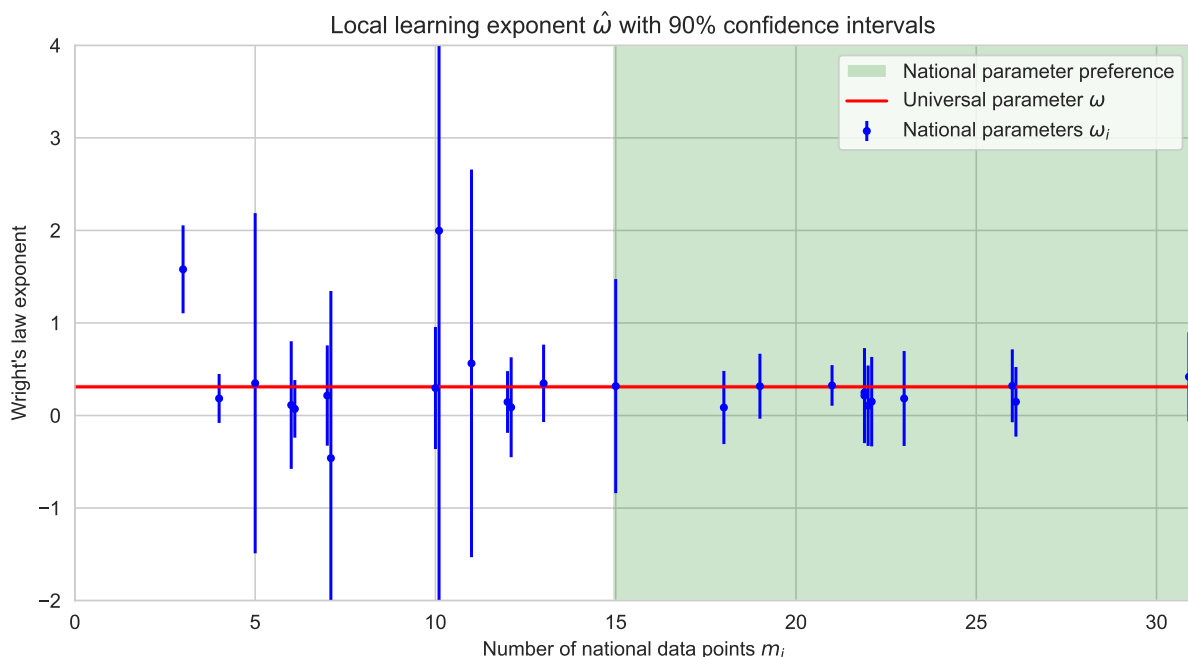


Figure B.11: **National Wright’s law parameters  $\omega_j$  against the number of available data points.** We show the estimated national Wright’s law exponents  $\hat{\omega}_j$  depending on the number of national data points available ( $m_j$ ). Blue dots indicate the parameters with the associated 90% confidence intervals. The confidence intervals are not symmetric due to a logarithmic y-scale. The red line shows the universal Wright’s law parameter  $\hat{\omega}$  from the *Wright, Local (U)* model. The green shaded area indicates parameters with  $m \geq 15$ . For these parameters, a national model would be preferred. We see that for the parameters in the green shaded area, the universal parameter  $\hat{\omega}$  always falls into the individual parameter’s 90% confidence intervals of  $\hat{\omega}_j$ .

exercise, we show that the uncertainty associated with the *Wright, Local (U)* model is larger than what we observe in the data. The reason for this lies in the second empirical argument. Unlike the *Wright, Local (U)* implies an increase in the cross-sectional variance that we do not observe in the real data.

Empirically, as we have seen in Figure B.10, there is no statistically significant difference in point-forecast accuracy between the *Wright, Local (U)*, *Moore, Local (U)*, and *Wright, Global* model. This is not surprising. We expect the *Wright, Local (U)* and *Moore, Local (U)* models to perform similarly a priori due to the exponential growth in cumulative national solar deployment. We also expect the *Wright, Local (U)* and *Wright, Global* models to perform similarly due to the high colinearity between national and global solar deployment.

To choose between models, we turn to the stochastic representation of out-of-sample

residuals. The model used to forecast residuals in the *Wright, Global* model is described in the main text and Section B.3.4. We apply the random walk model of Lafond et al. (2018) to forecast the global average and use an AR(1) model to describe the national deviations from that global average. For the *Wright, Local (U)* model, we consider the random walk representation from Lafond et al. (2018) on the national BOS costs:

$$\log \mathcal{J}_j(t) = \log \mathcal{J}_j(t-1) - \hat{\omega} \log \frac{Z_j(t)}{Z_j(t-1)} + \eta_{jt} + \rho \eta_{jt-1}, \quad (\text{B.22})$$

where  $\eta_{jt} \sim \mathcal{N}(0, \hat{\sigma}^2)$  i.i.d. and  $\hat{\omega}$  follows a student distribution around the true value of  $\omega$ .  $\hat{\omega}$  is estimated from a panel regression with fixed effects.  $\hat{\sigma}^2$  is estimated on the first log-differences of the cost and cumulative deployment.

Our previous scoring rule, the RMSE, does not capture the full forecast accuracy since it only considers the (median) point forecast. An alternative scoring rule that considers the full distribution of stochastic forecasts is the *Continuous Ranked Probability Score* (CRPS). The CRPS is commonly used to compare stochastic forecasts on continuous random variables. It is defined as

$$CRPS(F, y) = \int_{\mathbb{R}} (F(x) - \mathbb{I}_{(x \leq y)})^2 dx, \quad (\text{B.23})$$

where  $F(\cdot)$  is the cumulative density function of our forecast, and  $y$  is the empirical observation. If our forecast is a delta function that matches the true data exactly, i.e., we produce perfect forecasts and are certain that they are correct, then  $CRPS = 0$ . If either the forecast is off, or we are unsure how accurate the forecast is (or both), then  $CRPS > 0$ . In our case, we want to estimate the CRPS based on an ensemble of discrete stochastic point forecasts such that we can use Monte Carlo methods to approximate the CRPS.

There are different discrete estimators. In our case, we apply the so-called "energy

formulation“ CRPS:

$$\widehat{CRPS}(\{x_i\}_{i=1,\dots,M}, y) = \frac{1}{M} \sum_{i=1}^M |x_i - y| - \frac{1}{2M^2} \sum_{i,j=1}^M |x_i - x_j|, \quad (\text{B.24})$$

where  $\{x_i\}_{i=1,\dots,M}$  is the stochastic forecast ensemble and  $y$  is the true data we test against (Zamo and Naveau, 2018). We choose  $M = 500$  in the estimation below. It is non-trivial to show that Equation (B.23) and (B.24) are equivalent. The first term in (B.24) is the mean absolute error. The second term is independent of the true data  $y$  and, instead, describes the width of the forecast distribution. A narrow and accurate forecast will have low MSE and low forecast width. A narrow yet inaccurate forecast may have low width yet high MSE. A wider forecast will have a large width and high MSE, such that the CRPS quantifies the trade-off between having a very uncertain forecast (but captures the true value) to one that is less uncertain but may not capture the true value. While this estimator (B.24) is biased, it outperforms unbiased alternatives for large values of  $M$  and few observations  $y$  (Zamo and Naveau, 2018).

As for the RMSE, we compare the average CRPS of our two models for different forecast horizons  $\tau$ . As explained earlier, we focus on the forecast residual relative to the global average cost,  $\log \bar{\mathcal{J}}_j$  defined in Equation (B.17), since this significantly increases the statistical significance of our comparison.

The results of this analysis are shown in Figure B.12. Except for  $\tau = 1$ , the *Wright, Global (U)* model shows a lower CRPS than the *Wright, Local (U)* model. The average CRPS is smaller, and there are several individual data points for which the forecasts are more accurate under the *Wright, Global (U)* model. Since we do not find a different RMSE of both forecasts, this means that while both point forecasts have similar accuracy, the *Wright, Global* model has a smaller confidence interval around these forecasts. The *Wright, Local (U)* model over-estimates the future uncertainty around the point forecast.

The reason for this difference between the models lies in their forecast for cross-sectional averages. The *Wright, Global (U)* model predicts the deviation from the global average

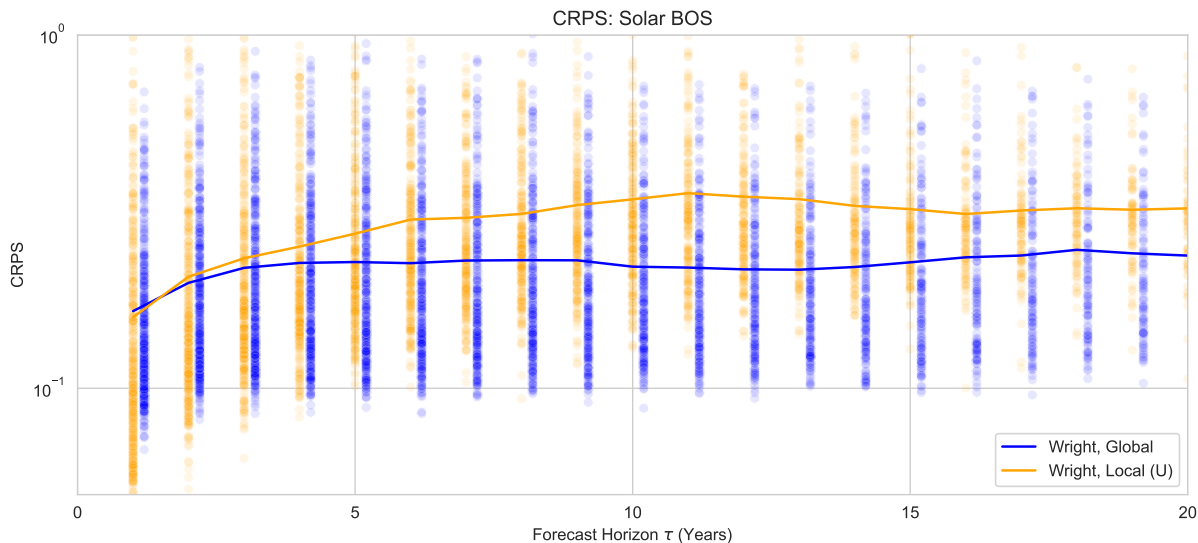


Figure B.12: **Out-of-sample comparison of the solar BOS cost forecast errors under the *Wright, Local (U)* and *Wright, Global (U)* model relative to the global average solar BOS costs, in terms of CRPS.** We consider only the national solar BOS cost relative to the global average BOS cost. We estimate model parameters on 5-year moving average windows on our data and forecast the costs relative to the true global average using the *Wright, Global (U)* (orange) and *Wright, Local (U)* model (blue). We make stochastic forecasts for the future of this window with a Monte Carlo approach. For each forecast, we estimate the CRPS using the energy formulation. The dots represent individual forecasts, and the lines are the average overall forecasts with each model. We see that, with the exception of  $\tau = 1$ , the *Wright, Global (U)* model outperforms the *Wright, Local (U)* model.

BOS costs with an AR(1) model. This means that the cross-sectional variance remains approximately constant over time. The *Wright, Local (U)* model, on the other side, predicts an increase in cross-sectional variance over time. Differences in national solar deployment infer different expected costs. Additionally, the stochastic noise is independent for each country, resulting in further divergence of national costs. In reality, we have seen in B.8 that the cross-sectional variance has declined historically. Since this is contrary to what we expect under the *Wright, Local (U)* model, we reject it and work with the *Wright, Global (U)* model going forward.

### B.3.5 Wind BOS cost forecasts

**Summary:** *We do not find significant historical cost declines in national wind BOS costs. Instead of an experience curve model, we forecast future global BOS costs as a mean-reverting process. Equivalent to solar, we forecast national deviations from the global*

	Solar PV			Wind		
Local learning $\omega$	-0.3114***		-0.1692***	-0.0218*		-0.0328
Global learning $\tilde{\omega}$		-0.3666***	-0.1785***		-0.0255	0.0175
Entity Effects	X	X	X	X	X	X
$R^2$	0.7895	0.7912	0.8151	0.0059	0.0044	0.0065

\*  $p < 0.1$ ; \*\*  $p < 0.05$ ; \*\*\*  $p < 0.01$

Table B.4: Regression results for BOS costs, according to Equation (B.25). For Solar, all regression parameters are highly significant. For wind, the local learning parameter is significant below 10% if we include only local learning.

*mean as a mean-reverting process.*

### In-sample analysis

Before developing our forecast model, we perform an in-sample analysis of the wind BOS costs. This demonstrates that wind BOS costs follow a constant trend and are significantly cross-correlated. These findings are robust with respect to the lag we use to calculate national BOS costs.

**Trend analysis** We perform a regression analysis on national BOS costs using the model

$$\log \mathcal{J}_j(t) = \alpha_j + \omega * \log Z_j(t) + \tilde{\omega} * \log Z(t). \quad (\text{B.25})$$

$\alpha_j$  is a country fixed effect,  $\omega$  is a Wright's law type learning parameter with respect to the local cumulative deployment  $\log Z_j(t)$ , and  $\tilde{\omega}$  is a Wright's law parameter with respect to the global cumulative deployment  $\log Z(t)$ .

The results of this analysis (and derivatives thereof) are shown in Table B.4. For solar, we see highly significant global and local learning. This is expected given the large cost decline and take up of solar electricity. For wind we do not see any significant global or local learning (except a 10% significance for local learning).

This absence of learning has been noted previously by [Key et al. \(2022\)](#) but has otherwise gotten little attention in the literature. In some sense, the BOS costs are comparable to price development in the fossil fuel sector, where prices have remained approximately

constant despite improvements in extraction technologies (Pindyck, 1999). As a result, we expect further learning effects in onshore wind to be substantially below those in solar. Turbine costs are already decreasing slower (with respect to their deployment) than solar modules and BOS costs, leading to significant short- and long-term differences in cost reduction.

**Cross-sectional correlation analysis** We apply the same cross-sectional correlation test to our data as we did for solar. Specifically, we observe significant cross-sectional correlation based on a cross-sectional dependence test from Pesaran (2021), see Equation (B.15). We need to take this into account in our forecast model.

**Autocorrelation analysis** We first investigate the autocorrelation in our global average BOS costs  $\log \mathcal{J}(t)$ . As for solar,  $\log \mathcal{J}(t)$  is calculated as the unweighted average across all  $\log \mathcal{J}_j(t)$ , where  $j$  is the country identifier and  $t$  is the year of observation. We then compute the autocorrelation function (ACF) and partial autocorrelation function (PACF) for the resulting time-series  $\log \mathcal{J}(t)$ . We plot the autocorrelation and partial autocorrelation function of the resulting time series in Figure B.13. The autocorrelation function (left panel) declines gradually. However, only the first two lags are significant, as indicated by the blue shading. The partial autocorrelation function (right panel) declines rapidly after the first significant lag. These functions indicate that the global average can be approximated by an AR(1) process.

We repeat a similar analysis for the national BOS costs. For each country in our sample, we normalize the time-series by subtracting the global average BOS costs, according to equation (B.17). We then calculate the ACF and PACF for  $\log \bar{\mathcal{J}}_j(t)$  in each country  $j$  and take the respective cross-sectional weighted averages. The weighting depends on the length of the national time-series since longer time-series have a lower measurement error for the ACF and PACF. The reason for taking the average is the high level of noise in individual countries. The cross-sectional average reduces that noise to provide reasonable autocorrelation functions. The resulting average functions are shown in the left and right




Figure B.13: **ACF and PACF for the cross-sectional average wind BOS costs.** We consider the cross-sectional (national) average BOS costs for wind. **(left panel)** The ACF plot for the resulting time series. **(right panel)** The partial PACF plot for the resulting time series. The ACF is declining gradually, although only the first three lags are significant, indicated by the blue shading. The PACF is declining more abruptly, with only the first lag being significant. These two plots indicate an AR process at the origin of our data.

panels of figure B.14. We find one significant lag in both the ACF or PACF. This means that we can approximate the  $\log \bar{\mathcal{J}}_j(t)$  time-series with a similar first-order autocorrelation process.

**Sensitivity analysis in BOS estimation** As described in the main text, we calculate the wind BOS costs using a 1.5-year lag for turbine costs, such that

$$\mathcal{J}_j(t) = \mathcal{I}_j(t) - \omega_{t-1}\mathcal{M}_j(t-1) - \omega_{t-2}\mathcal{M}_j(t-2), \quad (\text{B.26})$$

where  $\mathcal{I}(t)$  is the total investment cost and  $\mathcal{M}(t)$  the turbine costs in year  $t$ , and  $\omega_t = 1/2$  denotes the weights applied to individual years. While the length of this lag is based on previous literature, it is important to investigate how robust our results are against this assumption. We repeat the previous regression analysis of wind BOS costs for different values of  $\omega_t$  or, equivalently, different lags.

Table B.5 shows our previous regression results for different lags. We focus on the




Figure B.14: **Cross-sectional weighted average ACF and PACF for the national wind BOS cost deviations.** We consider the national BOS costs deviation from the global average. We calculate national ACF and PACF functions and take their cross-sectional averages. **(left panel)** The cross-sectional weighted average ACF plot. **(right panel)** The cross-sectional weighted average PACF plot. The weights are based on the length of each national time series. The line shows the weighted median ACF / PACF. The box shows the weighted interquartile range. The whiskers show the 90% confidence range. The dots show outliers exceeding this range. The ACF shows a marginally significant first lag. Similarly, the PACF shows a marginally significant first lag. These plots indicate a first-order autocorrelation.

1-dimensional regressions with respect to the global and local cumulative deployment. The results show that for different lags, there is no significant learning. All regression results are insignificant at 5%. This means that our regression results are robust with respect to the lag applied to the turbine costs and there is no significant learning effect.

Lag (Years)	0	0.5	1	1.5	2	2.5	3
Only Local learning $\omega$	-0.0112	-0.0164	-0.0196*	-0.0218*	-0.0216*	-0.0208*	-0.0145
Only Global learning $\tilde{\omega}$	-0.0006	0.0120	-0.0194	-0.0255	-0.0285*	-0.0310*	-0.0195
Entity Effects	X	X	X	X	X	X	X

\*  $p < 0.1$  ; \*\*  $p < 0.05$  ; \*\*\*  $p < 0.001$

Table B.5: Regression results for wind BOS costs, according to Equation (B.25), under different lags. We see that independent of the lag applied to the turbine costs, the cost improvement in the wind BOS costs is not significant.

## Wind BOS model validation

Based on these in-sample statistics, we develop a model that reflects the absence of learning effects and correlation structures. We take from the approach in [Way et al. \(2022\)](#) for fossil fuel prices and model future wind BOS costs as an AR(1) process around the historical average.

$$\log \mathcal{J}(t) = \phi \log \mathcal{J}(t-1) + \epsilon_t + \kappa, \quad (\text{B.27})$$

$$\log \mathcal{J}_j(t) - \log \mathcal{J}(t) \equiv \log \overline{\mathcal{J}}_j(t) = \varphi \log \overline{\mathcal{J}}_j(t-1) + \eta_{jt}. \quad (\text{B.28})$$

As for solar PV, Equation (B.27) describes the global average wind BOS cost  $\log \mathcal{J}(t)$ . Equation (B.28) describes the national deviation from the global average  $\log \overline{\mathcal{J}}_j(t)$ . Both the global average BOS costs (B.27), as well as the national deviations thereof (B.28) follow an AR(1) process. The global average is calculated as a cross-sectional average of national BOS costs in each year  $t$ . The national deviations are calculated as the difference from that average in each year  $t$ .

For both global average (B.27) and national deviations (B.28) we need to determine the autocorrelation parameters  $\phi$  and  $\varphi$ , and the size of the exogenous shocks  $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$  and  $\eta_{jt} \sim \mathcal{N}(0, \sigma_\eta^2)$ . Furthermore, the global average model (B.27) depends on the constant  $\kappa$ . For  $\phi$ , the short length of our time series makes it very difficult to estimate directly. Instead, we use the average value for fossil fuels in [Way et al. \(2022\)](#) of  $\phi = 0.88$ . For  $\varphi$ , the issue is similar. We, therefore, estimate  $\varphi_j$  based on a maximum likelihood estimation in all countries  $j$  and take the average thereof. This leaves us with  $\varphi \approx 0.7$ , similar to our model for solar BOS costs. For  $\hat{\kappa}$ , we use the long-term mean cost to estimate the constant

$$\hat{\kappa} = \frac{1}{m} \sum_t \log \mathcal{J}(t)(1 - \phi). \quad (\text{B.29})$$

The estimation error of  $\hat{\kappa}$  can be estimated based on the usual standard error in the mean. Finally, we estimate the variance of the i.i.d. shocks  $\epsilon$  and  $\eta$  based on the residuals of

Equations (B.27) and (B.28).

As before, statistical significance is an issue when it comes to validating this model due to the limited number of independent data points. The global model (B.27) is equivalent to the AR(1) model for fossil fuel costs in Way et al. (2022) and has been validated on a small set of long time series. We focus our validation on the national deviation model (B.28) where we can make statistically significant statements. We further do not include the estimation of  $\varphi$  in our validation but take it as a given input. Estimating the autocorrelation parameter on a short time series is challenging and limits our ability to validate the more general model and noise parameter  $\sigma_\eta$  (Farmer and Lafond, 2016).

Our procedure is the same as previously described; we consider 5-year moving average windows to estimate  $\sigma_\eta$  and make stochastic forecasts for the respective future of  $\log \bar{\mathcal{J}}_j(t)$ . We compare the distribution of our forecasts to that of the residuals across quantiles  $p$  using the cumulative PIT. We limit this analysis to forecasts made from 2003 onward. The reason for this is the limited data on national investment costs before that. Prior to 1999, we only have data for 6 countries, resulting in only 5 independent time series since one degree of freedom is used to estimate the global average BOS cost. Moreover, prominent data sources like IRENA (2024) show exactly equivalent costs for some countries, such as Germany, Sweden, Italy, Spain, and France, in 1993. It is unclear if the reason for this anomaly is due to issues with the underlying data or reflects an actual equivalence of costs between these countries.

Figure B.15 shows the observed cumulative PIT  $q_\Theta(p)$  against  $p$ . The grey bars are the empirical observations, and the red areas are the associated confidence intervals. As they were for the previous models, these are based on surrogate data. We see that the empirical quantiles are close to the identity line across all percentiles without any statistically significant differences.

### B.3.6 Investment cost forecasts

**Summary:** *We aggregate our forecasts for the modules/turbines and BOS costs to obtain a*




Figure B.15: **Out-of-sample validation of the national wind BOS cost forecast error, relative to the true global average cost.** We consider only the national deviation of wind BOS cost relative to the global average BOS cost (B.28). We estimate model parameters  $\sigma_\eta$  on 5-year moving average windows on our data and make stochastic forecasts for the relative costs using the AR(1) model described in the main text. The empirical cumulative PIT  $q_\Theta(p)$  is shown in grey. The red line is the theoretically expected PIT. The red areas are the 50% and 95% confidence intervals under a simulated AR(1) model. We see that the forecast distribution is statistically identical to the AR(1) model.

*forecast for the total investment cost. Due to the comparable learning rates of photovoltaic modules and BOS costs, our aggregate model performs similarly to a simpler Wright's law model on the investment cost. For wind, the difference in learning rates between turbines and BOS costs implies that our aggregate model is preferred over a pure Wright's law model.*

When we are interested in the total investment costs  $\mathcal{I}$ , we can sum the disaggregate module/turbine and BOS forecasts according to

$$\mathcal{I}_{jt} = \mathcal{M}_{jt} + \mathcal{J}_{jt}. \quad (\text{B.30})$$

Since the disaggregate model (based on  $\mathcal{M}$  and  $\mathcal{J}$ ) contains more information, we would expect it to always outperform alternative formulations on the aggregate  $\mathcal{I}$ . However, if we have little data to estimate our model parameters, using Wright's law on the total investment cost  $\mathcal{I}$  directly could be preferred over forecasting turbine/modules and BOS costs independently (Grunfeld and Griliches, 1960; Bermingham and D'Agostino, 2014; Hendry and Hubrich, 2011; Lütkepohl, 1984). The reason for this is the bias-variance

trade-off frequently encountered in forecasting (Baumgärtner et al., 2024). Furthermore, turbine, module, and BOS costs may be correlated with one another. In this case, the disaggregate model may overstate the future uncertainty on the total investment cost  $\mathcal{I}$ .

This topic warrants further study outside the scope of this work (Ferioli et al., 2009). Here, we perform two simple analyses to identify a model. Firstly, we investigate the cross-correlation between solar modules, wind turbines, and respective BOS costs. We then perform a heuristic test to identify which of the two investment cost model alternatives would perform better, conditional on the true underlying process. We find that the disaggregate model is more appropriate for wind and that both aggregate and disaggregate models can be used for solar. We have also performed an out-of-sample test comparing the forecast performance of different models using their RMSE. However, since the results are not significant we omit from reporting them here.

**Cross-correlation analysis** We consider the first difference of global average log costs for modules, turbines, and BOS costs. We then compute the correlation coefficient between the three time series using a Pearson test. We find that neither wind turbines and BOS nor solar modules and BOS costs are correlated with statistical significance. This indicates that we can use both forecast models independently and combine their results into an aggregate forecast.

**Simulation analysis** Using a simulation-based test, we show that the aggregate/disaggregate model choice most likely will not make a large difference for solar. For wind, the disaggregate model is likely the preferred choice.

We choose a simplified Moore’s law setup to simulate two time series  $X$  and  $Y$  with

$$X_t = X_{t-1}e^{\theta_X}e^{\epsilon_t} \tag{B.31}$$

$$Y_t = Y_{t-1}e^{\theta_Y}e^{\epsilon'_t}, \tag{B.32}$$

where  $\epsilon_t, \epsilon'_t \sim \mathcal{N}(0, \sigma^2)$  are i.i.d. noise shocks. We consider two settings. In the first,

we set both learning parameters  $\theta = -0.1$ , close to what we observe for solar modules and BOS costs. In the second, we consider  $X$  evolving as a geometric random walk with  $\theta = -0.1$  and  $Y$  evolving as an AR(1) process ( $\theta_Y = 0$  plus mean-reversion). In both settings, we simulate time series  $X$  and  $Y$  and fit both disaggregate and aggregate model parameters on a subset of  $m$  years. We use this to forecast 15 years into the future and calculate the Root Means Squared Error (RMSE) of these forecasts over a sample of 5000 simulation runs.

The results of this simulation are shown in Figure B.16. With equal learning, both models perform very similarly. The aggregate model outperforms the disaggregate for small training sets, but the difference is small and disappears if more than 25 data points are available. For the process with different learning rates, the two models perform differently. Even with few data points being used to estimate parameters, the disaggregate model outperforms the aggregate. This means that modellers may use either the aggregate or disaggregate model for solar but should use the disaggregate model in the case of wind.

### B.3.7 Capacity Factor forecasts

**Summary:** *We show that national capacity factors for solar and wind are approximately constant over time. We model capacity factors as i.i.d random variables. We compare two ways of estimating the capacity factor based on the annual electricity generation and the LCOE data. Using an out-of-sample test, we show that the LCOE-based estimation provides more accurate forecasts.*

We first investigate the change of national capacity factors over time. The literature disagrees on this topic, with some researchers depicting steep increases in capacity factor (IRENA, 2024; NREL, 2022), while others describe capacity factors as historically constant (Boretti and Castelletto, 2020; Bolinger and Seel, 2018; Boccard, 2009). We investigate trends in the capacity factor based on historical electricity generation and capacity.




Figure B.16: **Simulation of out-of-sample residuals for aggregate and disaggregate models, depending on the underlying processes.** We simulate the sum of two geometric random walks (**left**) and the sum of a geometric random walk and an AR(1) process (**right panel**). In each case, we apply a disaggregated (orange) and aggregated (blue) Moore’s law model with different sizes of the training set. We compare the out-of-sample residuals of both models for 15-year ahead forecasts using the RMSE. We see that for two geometric random walks, both models perform very similarly, particularly if more than 25 years of data are used to estimate the parameters. If one model evolves according to an AR(1) process, the disaggregate model significantly outperforms the aggregate, even with very few data being used to estimate model parameters.

### Capacity Factor stationarity

To assess the stationarity of a country’s capacity factor, we follow the approach by [Bolson et al. \(2022\)](#) and consider the actual national electricity generated per installed capacity. This approach differs from bottom-up modelled capacity factors, such as the one calculated by [ESMAP \(n.d.\)](#), which explicitly considers solar irradiation and other environmental factors. While the bottom-up approach is helpful for developing new solar sites and comparing technological choices, it can lead to significant estimation errors when compared with actual capacity factors ([Bolson et al., 2022](#)).

To identify trends in national capacity factors, we need to find a closed-form expression for the capacity factor of the newly built assets in a given time period. To do so, we must consider the large year-on-year growth in renewable capacity as well as the degradation  $D_j(t)$  of existing assets. Accrued over an asset’s lifetime, degradation of equipment has

a large effect on the capacity factor. Assuming a 20-year asset lifetime, solar equipment degrade by around 21% and wind assets by 28% (Staffell and Green, 2014; Kim et al., 2021).

Unfortunately, asset-level capacity and generation data are not widely available. Instead, our data is limited to the total annual electricity generation  $Y_j(t)$ , and the end-of-year electric capacity  $X_j(t)$ . One way to solve this constraint is to use a weighted average capacity factor  $\mathcal{Y}'_j(t)$  that we derive below.

For a given country  $j$ , we start by writing

$$Z_j(t) = Z_j(t - 1) + X_j(t), \quad (\text{B.33})$$

where  $Z_j(t)$  denotes the total national capacity at the end of year  $t$ , and  $X_j(t)$  is the additional capacity built throughout year  $t$ .

The total electricity generation  $Y_j(t)$  in year  $t$  depends on the installed capacity. Additionally, we need to account for the degradation  $D_j(t)$  of existing assets. We therefore write

$$\Delta Y_j(t) \equiv Y_j(t) - Y_j(t - 1) = \frac{1}{2} [\mathcal{Y}_j(t)X_j(t) + \mathcal{Y}_j(t - 1)X_j(t - 1)] - D_j(t). \quad (\text{B.34})$$

$\mathcal{Y}_j(t)X_j(t)$  is the average electricity generated by  $X_j(t)$ , dependent on the specific capacity factor  $\mathcal{Y}_j(t)$ . The factor of 1/2 accounts for the fact that  $X_j(t)$  is built throughout the year, assuming it is built halfway between each year. The degradation is taken as

$$D_j(t) = \delta * Y_j(t), \quad (\text{B.35})$$

where we assume a fixed annual decline in power output  $\delta$ . For wind, this degradation is around 1.6% annually (Staffell and Green, 2014), while for solar, it is around 1.2% annually (Kim et al., 2021). If we consider the weighted average capacity factor  $\mathcal{Y}'_t$  for

newly built resources in years  $t$  and  $t - 1$ , as

$$\mathcal{Y}'_j(t) = \frac{\mathcal{Y}_j(t)X_j(t) + \mathcal{Y}_j(t-1)X_j(t-1)}{X_j(t) + X_j(t-1)}. \quad (\text{B.36})$$

Instead of investigating  $\mathcal{Y}_j(t)$ , we consider  $\mathcal{Y}'_j(t)$  since it is mathematically easier, and any long-term trend in  $\mathcal{Y}_j(t)$  will translate into  $\mathcal{Y}'_j(t)$ . We write

$$\mathcal{Y}'_j(t) = \frac{2\Delta Y_j(t)}{X_j(t) + X_j(t-1)} + \frac{2D_j(t)}{X_j(t) + X_j(t-1)}. \quad (\text{B.37})$$

Based on formulation, we can now test for trend-stationarity in  $\mathcal{Y}'_j(t)$ . We assume a multiplicative lognormal error since the electricity generation can not be negative, and it is more natural to assume the error to scale with the deployed capacity. We also consider the autocorrelation between  $\mathcal{Y}'_j(t)$  and  $\mathcal{Y}'_j(t-1)$  since both values depend on the electricity produced in year  $t-1$ . This means that

$$\log \mathcal{Y}'_j(t) = \log [\mathcal{Y}_j(t)X_j(t) + \mathcal{Y}_j(t-1)X_j(t-1)] - \log [X_j(t) + X_j(t-1)] + \epsilon_t + \rho_t * \epsilon_{t-1}, \quad (\text{B.38})$$

where  $\epsilon_t \sim \mathcal{N}(0, \sigma_t^2)$  and  $\rho_t$  is the autocorrelation parameter. By the law of large numbers, this serves as a reasonable approximation of the mean 0. We can use this expression to consider the normalized difference between the logarithms of  $\mathcal{Y}'_j(t)$  and  $\mathcal{Y}'_j(t+\tau)$ . Under the null hypothesis, the difference is centred around 0, i.e.

$$H_0: \Delta_\tau \log \mathcal{Y}'_j(t) \equiv \frac{1}{\tau} \left( \log \mathcal{Y}'_j(t) - \log \mathcal{Y}'_j(t+\tau) \right) \sim \mathcal{N}\left(0, \frac{\sigma_t'^2 + \sigma_{t-2}'^2}{\tau^2}\right), \quad (\text{B.39})$$

where we discard the autocorrelation term  $\rho$  in place of  $\sigma'$ . For  $\tau = 1$ , this simply corresponds to taking the first difference in the logarithm of our adjusted capacity factor. Maximizing the data available to us, we consider all first differences for all countries. The distribution of  $\Delta_1 \log \mathcal{Y}'_j(t)$  (taking  $\tau = 1$ ) then behaves as a Gaussian mixture. Figure B.17 shows the observed distributions. Visually, we find that both distributions are symmetric




Figure B.17: **Observed distribution for the normalized log-difference of the capacity factor for newly installed capacity.** Blue bars show the distribution of observed values for  $\Delta_1 \log \mathcal{Y}'_j(t)$ . For missing country-year data points, we take the next best higher-order difference in  $\log \mathcal{Y}'_j(t)$  and normalize by the number of years included.

and centred around zero. Using the result of [Bakirov and Székely \(2006\)](#), we can apply a one-sided Student t-test to find that both for solar and wind, the deviation from zero is not significant at 90% confidence, with p-values of 61% for solar and 16% for wind. We note that this p-value provides an upper bound to the true statistical significance since we have not considered the likely high level of autocorrelation between the  $\log \mathcal{Y}'_j(t)$ .

This analysis shows that, while there has been some increase in capacity factor in some countries, we believe this to be due to a statistically insignificant global trend that will not continue in the future. For solar, this further means that global changes in the inverter loading ratio (ILR) ([IRENA, 2024](#); [Bolinger et al., 2022](#); [Deschamps and Rüther, 2019](#)) have not had a significant impact on the capacity factor. The caveat of this assumption is that the growing use of tracking modules, storage-integrated systems, as well as potential future use of bifacial and multi-junction modules may have some impact on the ILR and capacity factor ([SolarWorld, 2021](#); [de Souza Silva et al., 2021](#)).

## Capacity Factor model comparison

As we saw in the previous section, we can use the actual electricity generation and capacity to estimate the capacity factor. Alternatively, we can estimate the capacity factor from the LCOE data. In this section, we describe and compare both of these methods to identify the one best suited to our forecast application. Since we are using the electricity-based capacity factor to harmonize some of our wind LCOE data, the test is less meaningful than it is for solar. For completeness, we nevertheless report on both technologies here.

For the electricity model, we can use the average annual capacity factor (excluding outliers with an annual capacity factor of zero or exceeding 100% utilization). Specifically, we can use the unbiased estimator

$$\text{Electricity Model: } \hat{\mathcal{Y}}_j = \frac{1}{m} \sum_{t=1}^m \frac{Y_j(t)}{Z_j(t)}, \quad (\text{B.40})$$

where  $m$  is the number of observations.<sup>1</sup>

For the LCOE-based model, we can use Equation (B.11). We have seen in Section B.3.1 that  $\log \mathcal{E}$  follows an i.i.d. symmetric distribution. If we combine  $\log \mathcal{E}$  and  $\log \mathcal{Y}$  into a single random variable we can also estimate the capacity factor as

$$\text{Fixed-effect Model: } \log \hat{\mathcal{Y}}_j = \frac{1}{T_j} \sum_t (\log (\mathcal{I}_j(t) * \mathcal{C}_j(t) + O_j(t)) - \log \mathcal{L}_j(t)) \quad (\text{B.41})$$

where we sum over the  $T_j^k$  available data points for country  $j$  and technology  $k$ .

To choose between the electricity and fixed-effect models, we compare out-of-sample forecast errors in both models using backtesting. Since we are ultimately interested in forecasting the LCOE we compare our capacity-factor forecasts to the capacity factors observed in the LCOE:

$$\text{Observed data: } \log \mathcal{Y}_j(t) = \log (\mathcal{I}_j(t) * \mathcal{C}_j(t) + O_j(t)) - \log \mathcal{L}_j(t). \quad (\text{B.42})$$

---

<sup>1</sup>Unlike before, we no longer need to worry about correcting for capacity growth and asset deterioration since we have already shown that capacity factors are approximately constant over time.




Figure B.18: **Out-of-sample forecasts error for the capacity factor based on electricity generation and capacity, and a fixed-effect model for solar and wind.** We compare the out-of-sample forecast errors for the capacity factor in solar and wind. The x-axis shows the number of training years used in a moving window approach. The y-axis shows the error of the resulting forecasts. The box indicates the IQR, the whiskers extend this to 1.5 times the IQR. The gray boxes represent the electricity-based estimation, and the blue boxes represent the fixed effects estimation. We see that the fixed-effects forecasts outperform the electricity-based forecasts in both technologies and all training sizes.

We consider different training data in the form of moving average windows of varying sizes to estimate the capacity factor  $\hat{\mathcal{Y}}_j$  for wind and solar based on the two models.

The results are displayed in Figure B.18. We first notice that for both models, the error distribution is centred around zero. We also see that beyond 10 training data points, the errors do not decrease significantly as the training size increases. For solar, this means that we can estimate the capacity factor well, even if there are only a few data points available. Most importantly, the error is smaller for the fixed-effects model than it is for the electricity-based model. We, therefore, use the fixed-effects model in both technologies. For wind, this result is not meaningful, and we only include it in Figure B.18 for completeness. The reason for this is our data correction procedure, where we use the electricity-based capacity factor already to create a more consistent dataset.

### B.3.8 Operating & Maintenance cost forecasts

**Summary:** *Based on the limited data available, we show that operating & maintenance*




Figure B.19: **Annual operating costs for wind and solar electricity vs investment costs.** Total investment costs are plotted against annual operating costs per kW capacity for solar (**left**) and wind (**right**). The dotted lines indicate a fixed 3% relationship for wind and 1% for solar. Solar displays a Pearson correlation coefficient of  $\sim 0.58$ , wind  $\sim 0.67$  (Riva et al., 2018; Bolinger et al., 2022; Wiser et al., 2019).

*costs are approximately constant relative to the total investment costs. This means that we can use a constant to forecast the relative operating & maintenance costs.*

Figure B.19 shows O&M costs against the investment costs. We see that the observations are distributed symmetrically around the 3% O&M share for wind and 1% O&M share for solar. There is no clear differentiation between countries. Investment and O&M costs correlate positively with Pearson correlation coefficients of  $\sim 0.58$  for solar and  $\sim 0.67$  for wind.

### Stationarity relative to total investment cost

We hypothesize that the relative O&M cost  $\gamma_j = O_j/I_j$  is approximately constant across time and countries  $j$ . Since both investment and O&M costs differ across countries, we consider the average O&M cost in each country and year. This gives us a national time series of relative O&M costs. For each time series, we compute the first differences  $\Delta\gamma_j = \gamma_j(t) - \gamma_j(t-1)$ .

We confirm that there is no significant trend in  $\gamma_j$  by applying a two-sided t-test.

At p-values of 66% for wind and 9.4% for solar, neither show significant trends at 95% confidence. If we do not differentiate by country  $j$  and consider the first differences of the “global” average relative O& M cost in each year, these results do not change for wind. For solar, however, the time-trend is now significant at 95% confidence. As evident from Figure B.19, this result is driven primarily by the difference in national relative O& M cost and the fact that we only have data for two countries represented unequally in our sample.

For our forecasting purposes, we do not model a trend in relative O& M costs. Note that this result is consistent with the learning rates observed for O& M costs (Steffen et al., 2020).

### Operating & Maintenance cost model validation

Next, we estimate the average relative O&M share  $\log \hat{\gamma}$  and associated variance  $\hat{\sigma}$  from the data. We assume that the logarithm of the relative O&M share follows a normal distribution:

$$\log \frac{O_{ij}}{\mathcal{I}_{ij}} \sim \mathcal{N} \left( \log \hat{\gamma}, \hat{\sigma}^2 \left( \frac{1}{m} + 1 \right) \right), \quad (\text{B.43})$$

where  $m$  is the number of observations.

Based on this, we can estimate  $\log \hat{\gamma}$  as the sample mean and  $\hat{\sigma}$  as the sample variance of  $\log \frac{O_{ij}}{\mathcal{I}_{ij}}$ :

$$\log \hat{\gamma} = \frac{1}{m} \sum_{i=1}^m \log \frac{O_{ij}}{\mathcal{I}_{ij}}, \quad (\text{B.44})$$

$$\frac{\hat{\sigma}^2}{m^k} = \frac{1}{m} \frac{1}{m-1} \sum_{i=1}^m \left( \log \frac{O_{ij}}{\mathcal{I}_{ij}} - \log \hat{\gamma} \right)^2. \quad (\text{B.45})$$

Our stochastic model of  $\hat{\gamma}$  does not include the asset-specific exogenous shocks  $\epsilon_i$ . Since we are modelling the national weighted average relative O&M costs, these shocks are balanced out by the law of large numbers.

We validate this approach out-of-sample using a cross-sectional bootstrapping approach. Each country-year combination constitutes one data point. We repeatedly split this dataset

randomly into disjoint training and testing data, with  $m = 10$  data points in the training set and the remainder in the test set. We use the training data to estimate the model parameters  $\hat{\gamma}$  and  $\hat{\sigma}$ . For each observation in our testing data, we calculate the model error:

$$\mathcal{E}_i = \log \frac{O_{ij}}{\mathcal{I}_{ij}} - \log \hat{\gamma}. \quad (\text{B.46})$$

Since our training-samples will have higher or lower variance  $\hat{\sigma}$ , we normalize the error by the estimated standard deviation

$$\frac{1}{\sqrt{\frac{1}{m} + 1}} \frac{\mathcal{E}_i}{\hat{\sigma}} \sim t(m - 1). \quad (\text{B.47})$$

We expect the resulting normalized error distribution to follow a Student t-distribution with  $m - 1$  degrees of freedom.

The result of this analysis is shown in Figure B.20. We see that the t-distribution describes the out-of-sample error well and, thus, serves as an appropriate stochastic error model.

### B.3.9 Additional Cross-checks

In this section, we outline additional cross-checks performed on our data. Due to the correlation structures, these are not statistically significant but give some indication on potential overfitting by our models.

#### LCOE model

We compare the out-of-sample performance of our disaggregate LCOE model with two model alternatives: the *Wright, Global*, and *Wright, Local (U)* model. The *Wright, Global* model is based on Lafond et al. (2018) and follows the global Wright’s law. The *Wright, Local (U)* model is also a Wright’s law model. The difference to the *Wright, Global* model is that it forecasts future costs conditional on the national cumulative deployment of a

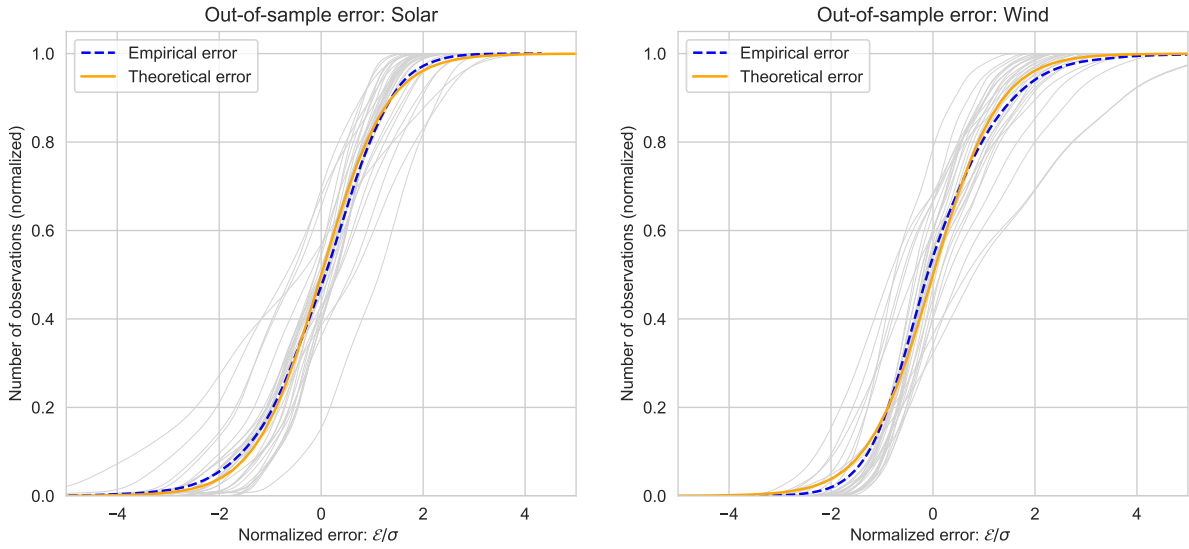


Figure B.20: **Observed out of sample error of the relative O&M share  $\log \gamma$  compared to the theoretical error.** The gray lines show the normalized out-of-sample errors of the observed relative O&M costs  $\log \gamma$ . The blue line shows the average over the gray lines. The orange line shows the theoretical distribution of the out-of-sample error, as per Equation (B.3.8). According to the error model, we expect the normalized out-of-sample error to follow a Student-t distribution. We observe a close match between the model and empirical data.

technology, not the global cumulative deployment. Since we will compare the same models for the solar BOS costs, we provide more detail in Table B.3.

We perform an out-of-sample test based on backtesting. We estimate the model parameters of all models on all 5-year moving windows, i.e., between times  $t$  and  $t - 4$ . We use these estimates to forecast future LCOE values in times  $t + \tau \geq t$ . To compare forecast residuals, we use the RMSE of Equation B.21.

The results of this analysis are shown in Figure B.21. The x-axis shows the forecast horizon  $\tau$ , and the y-axis shows the respective RMSE. Blue and purple lines represent the RMSE of the *Wright, Global*, and *Wright, Local (U)* forecasts. The black lines represent the RMSE of the disaggregate model. We see that the forecast accuracy of all models is similar. In fact, due to the small sample size, any observed difference is not statistically significant.




Figure B.21: **Out-of-sample residuals for different forecasting models of solar PV and wind LCOE.** We consider the out-of-sample forecasting error for solar PV (**top panel**) and wind (**bottom panel**) investment costs under different models described in Table B.3 and the previous sections. We estimate the model parameters on 5-year moving average training data for the forecast into the future of that respective training data and compare the forecasts with the actual data. The x-axis shows the forecast horizon. The y-axis shows the RMSE of the logarithm of the investment costs for our testing samples. All models perform similarly across forecasting horizons without statistically significant differences.

## Turbine cost model

In this section, we compare different turbine cost model alternatives. In particular, we compare if a model that differentiates different manufacturing locations is more accurate than a global model. Unfortunately, as we will show, there is not enough data to make a statistically robust statement.

We focus on three model alternatives described in Table B.6. We distinguish between Moore’s law models, which show cost decline over time, and Wright’s law models, which show cost declines conditional on cumulative deployment. For the Wright’s law model, we can condition future costs on local or global cumulative deployment. The letter (U) refers to the fact that we are using a universal model parameter  $\theta, \omega$  that is independent of the location  $j$ .

When comparing these three models, we focus on their out-of-sample accuracy in a backtesting exercise, as we have for the LCOE models. We consider all 5-year moving

Name	Description	Functional Form	Learning implications
Moore, Global (U)	We forecast $\mathcal{M}_j$ based on Moore’s law. The Moore’s law parameter $\theta$ is estimated using linear regression on the global average.	$\log \mathcal{M}_j(t) = \log \mathcal{M}_j(t-1) - \theta$	Both markets learn over time at roughly equivalent rates. The future costs in each market are independent of wind deployment.
Wright, Global (U)	We forecast $\mathcal{M}_j$ based on Wright’s law, conditional on the global cumulative deployment $Z(t)$ . The learning parameter $\omega$ is estimated on the global average, as in Way et al.way22.	$\log \mathcal{M}_j(t) = \log \mathcal{M}_j(t-1) - \omega \log \frac{Z(t)}{Z(t-1)}$	Both markets learn with increasing global experience. There is significant learning spillover through shared experience, i.e., markets depend on each others deployment. As learning effects dampen, these effects become less pronounced.
Wright, Local (U)	We forecast $\mathcal{M}_j$ based on Wright’s law, conditional on the local cumulative deployment $Z_j(t)$ . We use a universal Wright’s law exponent $\omega$ for both markets estimated with a panel regression.	$\log \mathcal{M}_j(t) = \log \mathcal{M}_j(t-1) - \omega \frac{Z_j(t)}{Z_j(t-1)}$	Both markets learn with increasing local experience. There is no spillover between markets and costs are independent of deployment in the other market. Costs are likely to diverge if both markets’ cumulative deployment grows at different rates.

Table B.6: **Overview of compared forecast models for wind turbine costs.** We compare 3 different forecast models forecasting future turbine costs manufactured in the West and China.  $\theta$  and  $\omega$  are Moore’s and Wright’s law parameters, respectively, independent of the market. We compare the model accuracy empirically to find that not one model outperforms the others with statistical significance.

time windows between years  $t$  and  $t - 4$  to estimate the model parameters ( $\theta$  and  $\omega$ ). We make forecasts for years  $t + \tau \geq t$  and compute the out-of-sample residuals for all forecasts with respect to the true data. We choose to evaluate forecast accuracy on the RMSE since we expect the forecast error to be distributed approximately lognormal based on typical Moore’s / Wright’s law error distributions (Farmer and Lafond, 2016; Lafond et al., 2018). In this case, the RMSE equals the sample standard error.

The result of this analysis is shown in Figure B.22. For all models, the RMSE increases with  $\tau$ . This is to be expected since we know less about the far future than we do about the near future. All three models perform very similarly. The *Wright, Global (U)* model appears marginally better than the other models. However, based on a Barlett’s test for each value of  $\tau$ , this difference is not statistically significant. The test is not conclusive, and based on the available data, we do not know which of the models provides the best




Figure B.22: **Out-of-sample comparison of different wind turbine forecast models.** We estimate model parameters on 5-year moving windows of our data. We make point forecasts for the future of this window and compare the forecasts to actual data. The x-axis shows the number of years we are forecasting into the future (forecast horizon). The y-axis shows the forecasts’ RMSE. We calculate the RMSE on the log and plot the relative error on the Turbine costs; an RMSE of 2 implies that, on average, our forecasts are off by a factor of 2. All three models show very similar RMSE for all forecast horizons  $\tau$ ; any difference in RMSE is not statistically significant.

forecasts.

For the Moore’s vs. Wright’s law models, this result is not surprising. As noted in earlier publications, the reason for the comparable predictive power in both models is the exponentially growing cumulative deployment of a technology (Lafond et al., 2020). Similarly, the cumulative deployment in each individual market is also highly colinear despite the ramp-up of Chinese turbine manufacturing, as confirmed by a Pearson test. In the presence of strong colinearity, the global and local models will also be equivalent. We thus cannot reject any of these three models based on the data alone.

## B.4 IAM scenario comparison

In this section, we provide more detail to the IAM scenario comparison presented in the main methods section. We first describe the data collection from the AR6 database (Byers et al., 2022). In Section B.6, we show additional country examples. In Section B.4.2, we show that if scenarios are corrected for the initial value, the bias in their forecasts is reduced substantially.

### B.4.1 Scenario data collection

We construct an ensemble of scenarios from the AR6 database (Byers et al., 2022). Since not all of these scenarios allow for a fair comparison, we exclude the following subset in our analysis:

- *Failed Vetting* These scenarios are identified by the AR6 authors; their input parameters were not considered “within reasonable ranges for the baseline period – primarily for indicators relating to emissions and the energy sector”
- *Negative growth* 5 scenarios for solar and 20 scenarios for wind show negative growth in the operational solar and wind capacity.
- *Negative cost* 5 scenarios report negative total investment costs for solar.

The resulting ensemble includes 741 scenarios for solar and 516 scenarios for wind. We use 2020 as the base year of our scenario comparison and account for inflation by scaling all cost data to 2020 USD (from 2010 USD as reported in the scenario database).

### B.4.2 Initial value correction

To understand the difference between our forecast results and IAM outputs we investigate the initial values used in both. In particular, the initial (2020) cost is an important input that can significantly change forecast results. This difference is the primary source of the observed bias in the scenarios.

Figure B.23 displays the reported solar and wind capacity for 2020 against the actual data. We see that in both cases, the reported investment cost is generally below what actually occurred. For solar, Remind 2.1 is the only model that reports costs below what was observed.

To understand the impact of this initial value, we correct scenario projections using a multiplicative factor and repeat the previous comparison between the total investment cost reported by the scenarios and the forecast by our model. The results of this for the

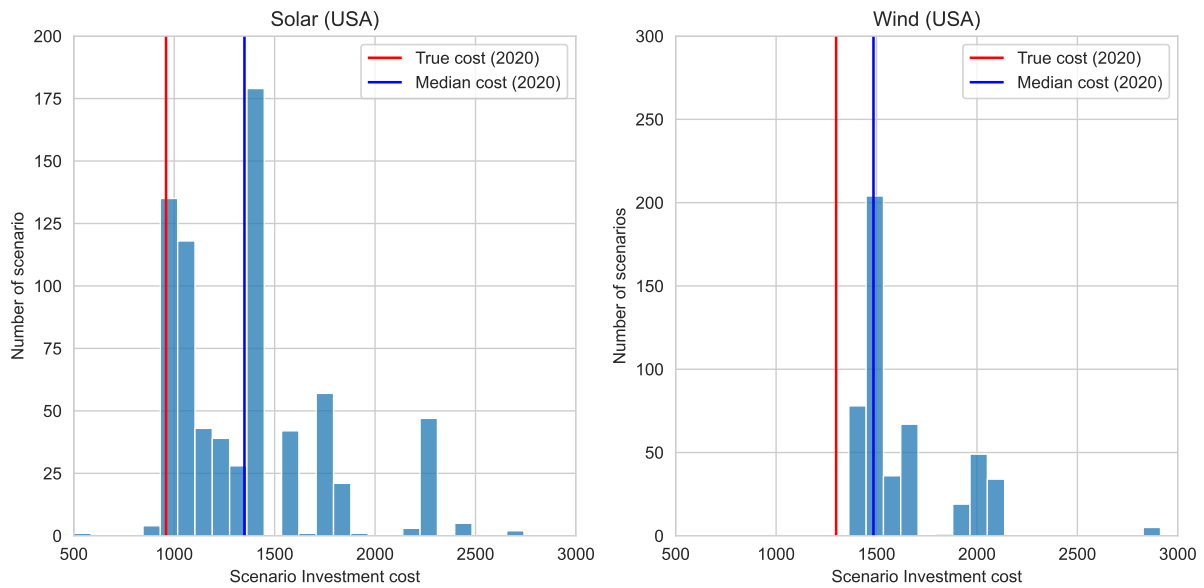


Figure B.23: **Current solar and wind total investment cost for the US as reported in IAM scenarios** The blue bars show the solar (left) and wind (right) total investment costs in 2020, as reported by the AR6 scenarios with regionally differentiated capital costs. We exclude scenarios with costs exceeding 3000 USD/kW for better visibility. The blue line indicates the median cost of each technology over the scenario sample. The red line shows the true average solar and wind cost in 2020. We see that almost all scenarios show costs that are larger than what was reported in 2020.

US are shown in Figure B.24. We see that for solar, the scenario costs now align much better with our median forecast. Now, only 25% of the scenarios exceed our point forecast. For wind, we observe that scenario costs are slightly below our point forecast but are still mostly within the 50% confidence interval.

One reason for this discrepancy in input data and subsequent differences in output scenarios is that many scenarios in the AR6 database were generated before 2020. For those scenarios, the 2020 cost and deployment are not based on actual data but on projections made in earlier years. In that case, even small biases in the early years of a scenario can lead to significant differences in the respective future cost projections. These challenges with IAMs have been noted previously, for example, [Creutzig et al. \(2017\)](#) and [Way et al. \(2022\)](#). Consequentially, our results emphasize the need for more stringent scenario validation and more frequent updates to the scenarios as the energy transition evolves.

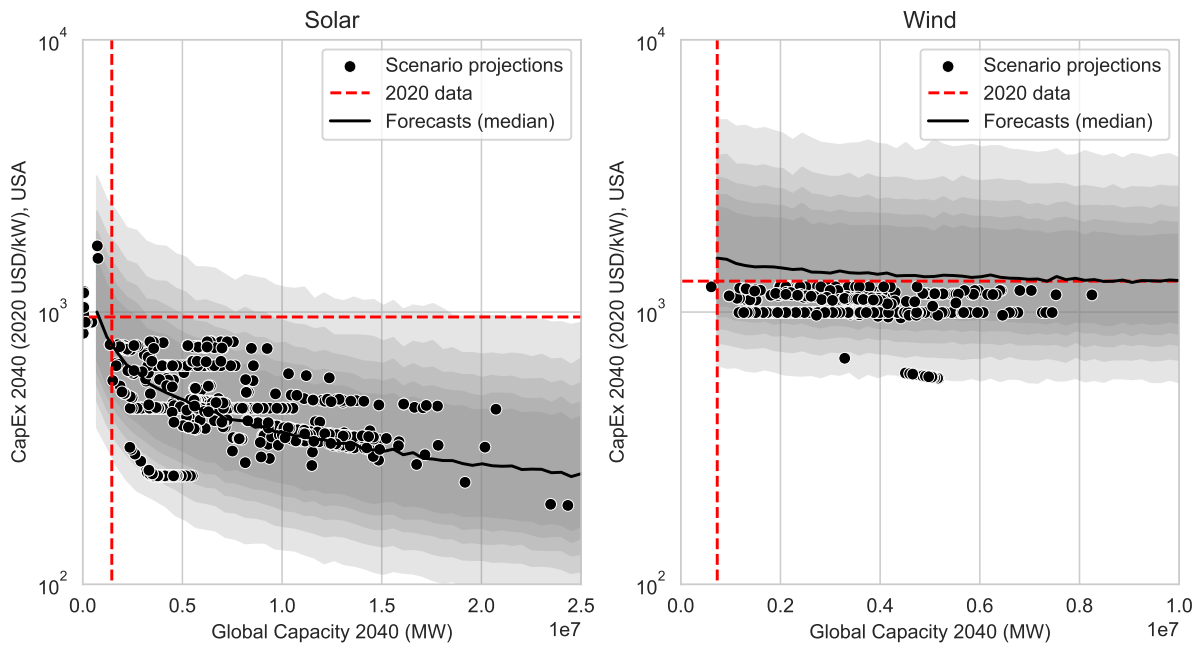


Figure B.24: **Adjusted Solar and wind total investment cost projections reported in IAM scenarios against probabilistic projections from our model.** The dots show the total investment cost projections for the US in different scenarios, adjusted for different initial costs. The x-axis shows the associated global solar / wind capacity. We compare them with the investment cost projections obtained from our model. The areas show 50-90% confidence intervals in 10% steps. The black line shows the median cost from our model. The red dotted lines show the global capacity and investment cost in the US for 2020. In both cases, the new investment cost lowers the expected future costs across almost all scenarios. We see that for solar, the scenarios now align well with our median forecast. For wind, the scenarios show costs that are slightly lower than what we expect. They are, however, still mostly within the 50% confidence interval.

## B.5 Estimated model parameters

The model parameters that were calibrated in this study are reported separately.

## B.6 Additional Figures

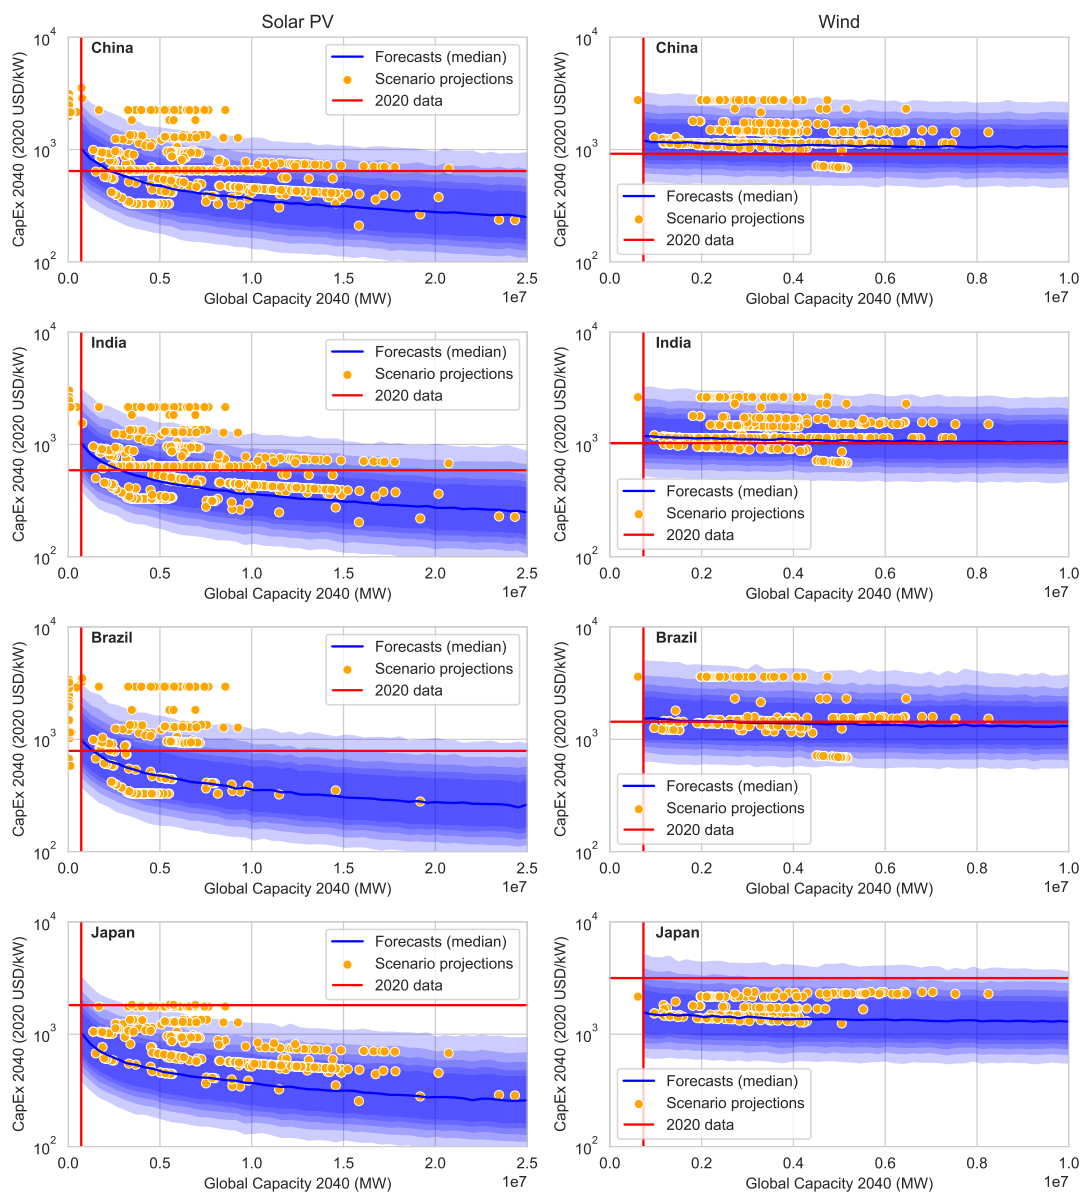


Figure B.25: **Solar (left) and wind (right) total investment cost projections reported in IAM scenarios against probabilistic projections from our model.** The orange dots show the total investment cost projections for the different countries in different scenarios. The x-axis shows the associated global solar / wind capacity. We compare them with the investment cost projections obtained from our model. The blue areas show 50-90% confidence intervals in 10% steps. The blue line shows the median cost from our model. The red lines show the global capacity and investment cost in each country for 2020. For wind in Brazil, the IAM projections generally align well with the median forecast. For the other countries and solar, we observe some discrepancies. In particular, many IAM scenarios are above the expected future costs, although they generally fall within the 50% confidence interval of our forecasts. In both cases, the probabilistic forecasts show significantly more uncertainty around the median forecasts than the IAM projections represent.

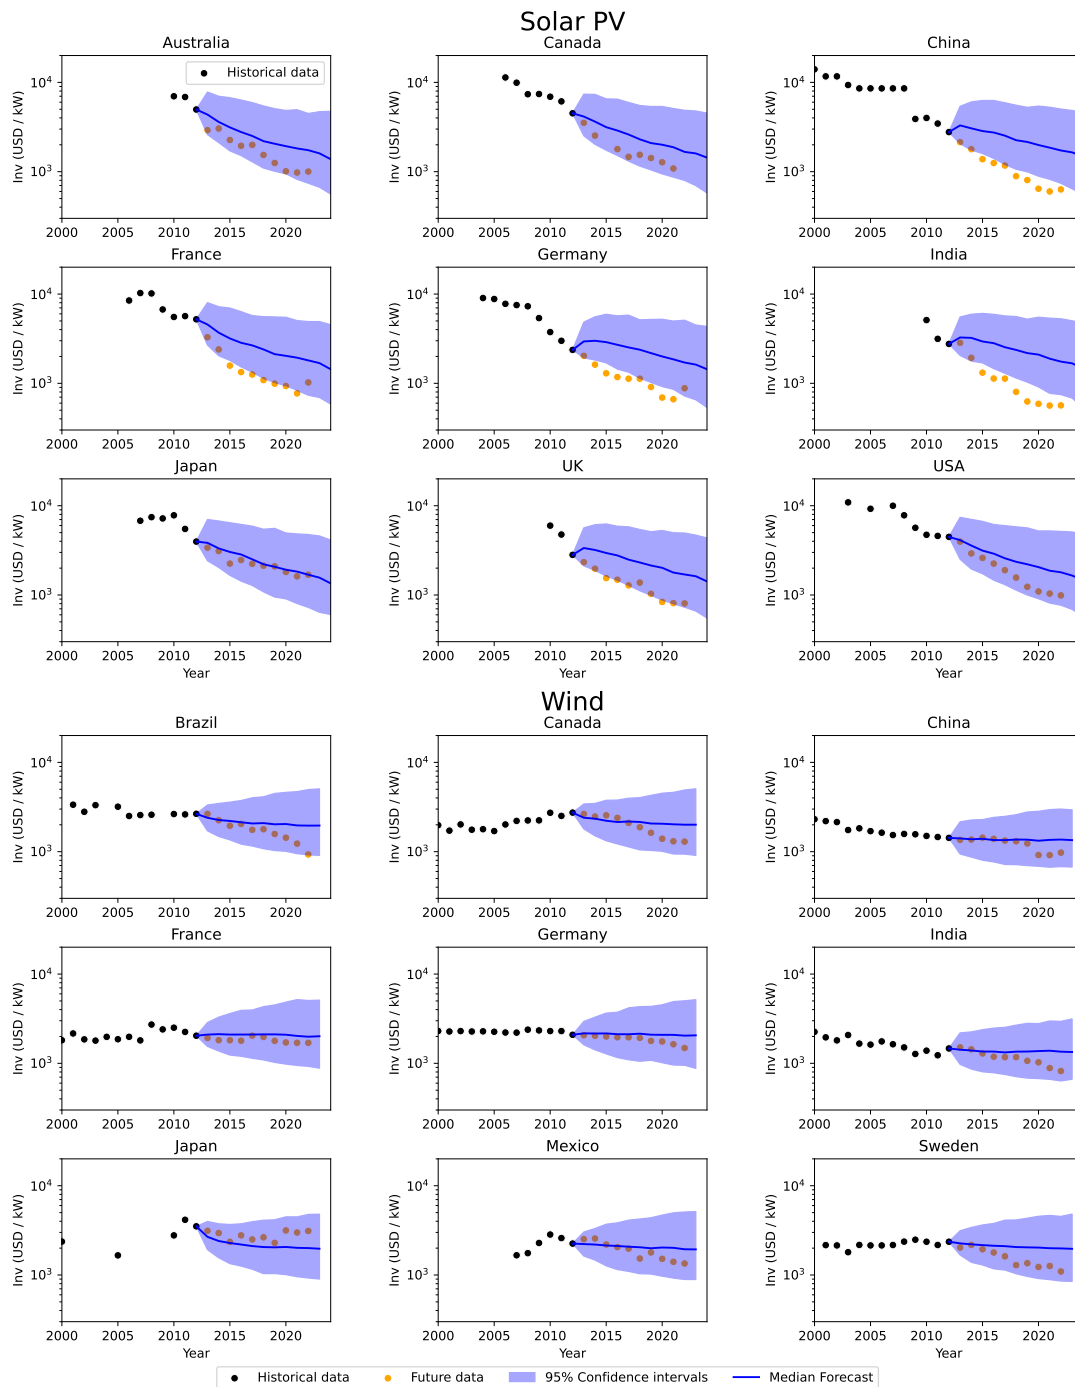


Figure B.26: **Investment cost forecasts compared to true data.** We illustrate the Investment forecast model by comparing the stochastic model outcomes to the actual data for selected countries (in-sample). The top 9 panels show national forecasts for solar, and the bottom 9 panels show national forecasts for wind. The y-axes show the national Investment cost on a logarithmic scale. The black dots indicate historical data, and the orange dots indicate future data. The blue shaded areas show the 95% confidence interval of our forecasts, and the blue line is the median forecast. With one exception in India, all future data is captured by the 95% confidence interval. For solar, we see that our forecasts are too high in all countries.

## References

- Bakirov, N. K. and G. J. Székely**, “Student’s t-Test for Gaussian Scale Mixtures,” *Journal of Mathematical Sciences*, December 2006, 139 (3), 6497–6505.
- Barbose, Galen, Naïm Richard Darghouth, Eric O’Shaughnessy, and Forrester**, “Tracking the Sun 2024 Edition,” Technical Report, Lawrence Berkley National Laboratory August 2024.
- Baumgärtner, C. Lennart, Rupert Way, Matthew C. Ives, and J. Doyne Farmer**, “The Need for Better Statistical Testing in Data-Driven Energy Technology Modeling,” *Joule*, September 2024, 8 (9), 2453–2466.
- Beattie, Shaun**, “How Long Does It Take To Develop, Design, And Deliver A Commercial Solar PV Installation?,” October 2021.
- Bermingham, Colin and Antonello D’Agostino**, “Understanding and Forecasting Aggregate and Disaggregate Price Dynamics,” *Empirical Economics*, March 2014, 46 (2), 765–788.
- Bloomberg**, “Bloomberg Terminal,” Bloomberg Finance L.P. 2022.
- Boccard, Nicolas**, “Capacity Factor of Wind Power Realized Values vs. Estimates,” *Energy Policy*, July 2009, 37 (7), 2679–2688.
- Bolinger, Mark and Joachim Seel**, “Utility-Scale Solar: Empirical Trends in Project Technology, Cost, Performance, and PPA Pricing in the United States – 2018 Edition,” Technical Report 1477381 September 2018.
- , – , **Cody Warner, and Dana Robson**, “Utility-Scale Solar, 2022 Edition: Empirical Trends in Deployment, Technology, Cost, Performance, PPA Pricing, and Value in the United States,” Technical Report None, 1888246, ark:/13030/qt7496x1pc September 2022.

**Bolson, Natanael, Pedro Prieto, and Tadeusz Patzek**, “Capacity Factors for Electrical Power Generation from Renewable and Nonrenewable Sources,” *Proceedings of the National Academy of Sciences*, December 2022, *119* (52), e2205429119.

**Boretti, Alberto and Stefania Castelletto**, “Trends in Performance Factors of Wind Energy Facilities,” *SN Applied Sciences*, September 2020, *2* (10), 1718.

**BP**, “Statistical Review of World Energy,” 2022.

**Bretherton, Christopher S., Martin Widmann, Valentin P. Dymnikov, John M. Wallace, and Ileana Bladé**, “The Effective Number of Spatial Degrees of Freedom of a Time-Varying Field,” *Journal of Climate*, July 1999, *12*, 1990–2009.

**Byers, Edward, Volker Krey, Elmar Kriegler, Keywan Riahi, Roberto Schaeffer, Jarmo Kikstra, Robin Lamboll, Zebedee Nicholls, Marit Sandstad, Chris Smith, van der Wijst, Kaj, Al -Khourdajie, Alaa, Franck Lecocq, Portugal-Pereira, Joana, Yamina Saheb, Anders Stromman, Harald Winkler, Cornelia Auer, Elina Brutschin, Matthew Gidden, Philip Hackstock, Mathijs Harmen, Daniel Huppmann, Peter Kolp, Claire Lepault, Jared Lewis, Giacomo Marangoni, Müller-Casseres, Eduardo, Ragnhild Skeie, Michaela Werning, Katherine Calvin, Piers Forster, Celine Guivarch, Tomoko Hasegawa, Malte Meinshausen, Glen Peters, Joeri Rogelj, Bjorn Samset, Julia Steinberger, Massimo Tavoni, and van Vuuren, Detlef**, “AR6 Scenarios Database,” November 2022.

**Comission de regulation de L’energie**, “Coûts et rentabilité des énergies renouvelables en France métropolitaine : Éolien terrestre, biomasse, solaire photovoltaïque,” April 2014.

**Creutzig, Felix, Peter Agoston, Jan Christoph Goldschmidt, Gunnar Luderer, Gregory Nemet, and Robert C. Pietzcker**, “The Underestimated Potential of Solar Energy to Mitigate Climate Change,” *Nature Energy*, August 2017, *2* (9).

- Cruce, Jesse, Eric O’Shaughnessy, Jenna Harmon, Jesse Geiger, and Jeffrey Cook**, “Residential Solar Adoption Timelines and Impacts from the COVID-19 Pandemic,” Technical Report NREL/TP-6A20-83529, 1889272, MainId:84302 September 2022.
- de Souza Silva, João Lucas, Karen Barbosa de Melo, Tatiane Silva Costa, Gilson Mario Vieira Machado, Hugo Soeiro Moreira, and Marcelo Gradella Villalva**, “Impact of Bifacial Modules on the Inverter Clipping in Distributed Generation Photovoltaic Systems in Brazil,” in “2021 Brazilian Power Electronics Conference (COBEP)” November 2021, pp. 1–6.
- Deschamps, Eduardo Martins and Ricardo Rüther**, “Optimization of Inverter Loading Ratio for Grid Connected Photovoltaic Systems,” *Solar Energy*, February 2019, *179*, 106–118.
- Egli, Florian, Bjarne Steffen, and Tobias S. Schmidt**, “A Dynamic Analysis of Financing Conditions for Renewable Energy Technologies,” *Nature Energy*, December 2018, *3* (12), 1084–1092.
- Elshurafa, Amro M., Shahad R. Albardi, Simona Bigerna, and Carlo Andrea Bollino**, “Estimating the Learning Curve of Solar PV Balance-of-System for over 20 Countries: Implications and Policy Recommendations,” *Journal of Cleaner Production*, September 2018, *196*, 122–134.
- ESMAP**, “Global Solar Atlas,” <https://globalsolaratlas.info>.
- Farmer, J. Doyne and François Lafond**, “How Predictable Is Technological Progress?,” *Research Policy*, April 2016, *45* (3), 647–665.
- Feroli, F., K. Schoots, and B. C.C. van der Zwaan**, “Use and Limitations of Learning Curves for Energy Technology Policy: A Component-Learning Hypothesis,” *Energy Policy*, July 2009, *37* (7), 2525–2535.

**Gneiting, Tilmann and Adrian E Raftery**, “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association*, March 2007, *102* (477), 359–378.

– **and Matthias Katzfuss**, “Probabilistic Forecasting,” *Annual Review of Statistics and Its Application*, January 2014, *1* (1), 125–151.

– **, Fadoua Balabdaoui, and Adrian E. Raftery**, “Probabilistic Forecasts, Calibration and Sharpness,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, April 2007, *69* (2), 243–268.

**Gosens, Jorrit and Yonglong Lu**, “Prospects for Global Market Expansion of China’s Wind Turbine Manufacturing Industry,” *Energy Policy*, April 2014, *67*, 301–318.

**Grunfeld, Yehuda and Zvi Griliches**, “Is Aggregation Necessarily Bad?,” *The Review of Economics and Statistics*, 1960, *42* (1), 1–13.

**Hayashi, Daisuke, Joern Huenteler, and Joanna I. Lewis**, “Gone with the Wind: A Learning Curve Analysis of China’s Wind Power Industry,” *Energy Policy*, September 2018, *120*, 38–51.

**Hendry, David F. and Kirstin Hubrich**, “Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate,” *Journal of Business & Economic Statistics*, 2011, *29* (2), 216–227.

**Hoyos, Rafael E. De and Vasilis Sarafidis**, “Testing for Cross-Sectional Dependence in Panel-Data Models,” *The Stata Journal*, November 2006, *6* (4), 482–496.

**IEA**, “World Energy Balances,” <https://doi.org/10.1787/data-00512-en>.

**IHSM**, “Electric Plant Unit Wind Turbine Installed,” March 2023.

**IMF**, “IMF Data,” <https://data.imf.org> 2022.

**IRENA**, “Renewable Power Generation Costs in 2021,” Technical Report, IRENA 2021.

- IRENA**, “Renewable Capacity Statistics 2022,” Technical Report, International Renewable Energy Agency, Abu Dhabi April 2022.
- IRENA**, “The Cost of Financing for Renewable Power,” Technical Report, International Renewable Energy Agency, Abu Dhabi 2023.
- IRENA**, “Renewable Power Generation Costs in 2023,” Technical Report, IRENA, Abu Dhabi 2024.
- Kempa, Karol, Ulf Moslener, and Oliver Schenker**, “The Cost of Debt of Renewable and Non-Renewable Energy Firms,” *Nature Energy*, February 2021, *6* (2), 135–142.
- Key, Alicia, Owen Robers, and Annika Eberle**, “Scaling Trends for Balance-of-System Costs at Land-Based Wind Power Plants: Opportunities for Innovations in Foundation and Erection,” *Wind Engineering*, 2022, *46* (3), 896–913.
- Kim, Jaeun, Matheus Rabelo, Siva Parvathi Padi, Hasnain Yousuf, Eun-Chel Cho, and Junsin Yi**, “A Review of the Degradation of Photovoltaic Modules for Life Expectancy,” *Energies*, January 2021, *14* (14), 4278.
- Kothari, Sumit, Nadia Ameli, and Michael Grubb**, “Drivers of Balance-of-System Costs for Solar PV Technology,” 2023.
- Lafond, François, Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J. Doyne Farmer**, “How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts,” *Technological Forecasting and Social Change*, March 2018, *128*, 104–117.
- , **Diana Seave Greenwald, and J. Doyne Farmer**, “Can Stimulating Demand Drive Costs Down? World War II as a Natural Experiment,” *SSRN Electronic Journal*, June 2020.
- Lütkepohl, Helmut**, “Forecasting Contemporaneously Aggregated Vector ARMA Processes,” *Journal of Business & Economic Statistics*, 1984, *2* (3), 201–214.

**Masson, Gaetan and Izumi Kaizuka**, “Trends in Photovoltaic Applications 2022,” Technical Report 43:2022, IEA PVPS TCP 2022.

**NREL**, “Annual Technology Baseline 2022,” <https://atb.nrel.gov/> 2022.

**Ondraczek, Janosch, Nadejda Komendantova, and Anthony Patt**, “WACC the Dog: The Effect of Financing Costs on the Levelized Cost of Solar PV Power,” May 2013.

**Pesaran, M. Hashem**, “General Diagnostic Tests for Cross-Sectional Dependence in Panels,” *Empirical Economics*, January 2021, *60* (1), 13–50.

– , **Andreas Pick, and Allan Timmermann**, “Forecasting with Panel Data: Estimation Uncertainty versus Parameter Heterogeneity,” April 2024.

**Pindyck, Robert S.**, “The Long-Run Evolution of Energy Prices,” *The Energy Journal*, 1999, *20* (2), 1–27.

**Polzin, Friedemann, Mark Sanders, Bjarne Steffen, Florian Egli, Tobias S. Schmidt, Panagiotis Karkatsoulis, Panagiotis Fragkos, and Leonidas Parousos**, “The Effect of Differentiating Costs of Capital by Country and Technology on the European Energy Transition,” *Climatic Change*, July 2021, *167* (1-2).

**PVPS, IEA**, “Trends in Photovoltaic Applications 2015,” Technical Report 27:2015, IEA PVPS TCP 2015.

**Raupach-Sumiya, Jörg, Hironao Matsubara, Andreas Prah, Astrid Aretz, and Steven Salecki**, “Regional Economic Effects of Renewable Energies - Comparing Germany and Japan,” *Energy, Sustainability and Society*, December 2015, *5* (1), 10.

**Reuters**, “Refinitiv Eikon,” Reuters 2022.

**Riva, Alberto Dalla, Janos Hethey, Silke Luers, Anna-Kathrin Wallasch, Knud Rehfeldt, Aidan Duffy, David Edward Weir, Maria Stenkvist, Andreas**

**Uihlein, Tyler Stehly, and Eric Lantz**, “IEA Wind TCP Task 26: Wind Technology, Cost, and Performance Trends in Denmark, Germany, Ireland, Norway, Sweden, the European Union, and the United States: 2008-2016,” Technical Report, National Renewable Energy Laboratory (NREL), Golden, CO (US) November 2018.

**Smart, Bridget and J. Doyne Farmer**, “Macro and Micro Prediction Models,” October 2024.

*Solar (Photovoltaic) Panel Prices*

**Solar (Photovoltaic) Panel Prices**, <https://ourworldindata.org/grapher/solar-pv-prices>.

**SolarWorld**, “Calculating the Additional Energy Yield of Bifacial Solar Modules,” Technical Report, SolarWorld September 2021.

**Staffell, Iain and Richard Green**, “How Does Wind Farm Performance Decline with Age?,” *Renewable Energy*, June 2014, 66, 775–786.

**Steffen, Bjarne**, “The Importance of Project Finance for Renewable Energy Projects,” *Energy Economics*, January 2018, 69, 280–294.

–, “Estimating the Cost of Capital for Renewable Energy Projects,” *Energy Economics*, May 2020, 88.

–, **Martin Beuse, Paul Tautorat, and Tobias S. Schmidt**, “Experience Curves for Operations and Maintenance Costs of Renewable Energy Technologies,” *Joule*, February 2020, 4 (2), 359–375.

**Thornton, A Grant**, “Africa Renewable Energy Discount Rate Survey – 2018,” 2018.

**U.S. Department of Energy**, “Land-Based Wind Energy Economic Development Guide,” Technical Report, Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office: WindExchange 2022.

**Way, Rupert, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer**, “*Empirically Grounded Technology Forecasts and the Energy Transition*,” *Joule*, September 2022, 6, 1–26.

**Wiser, Ryan, Mark Bolinger, and Eric Lantz**, “*Assessing Wind Power Operating Costs in the United States: Results from a Survey of Wind Industry Experts*,” *Renewable Energy Focus*, September 2019, 30, 46–57.

– , – , and **Joachim Seel**, “*Benchmarking Utility-Scale PV Operational Expenses and Project Lifetimes: Results from a Survey of U.S. Solar Industry Professionals*,” *Technical Report None*, 1631678, [ark:/13030/qt2pd8608q](https://doi.org/10.1039/c9te00088a) June 2020.

**Wood Mackenzie**, “*China Leads Global Wind Turbine Manufacturers’ Market Share in 2023*,” <https://www.woodmac.com/press-releases/2024-press-releases/global-wind-oem-marketshare/> May 2024.

**Zamo, Michaël and Philippe Naveau**, “*Estimation of the Continuous Ranked Probability Score with Limited Information and Applications to Ensemble Weather Forecasts*,” *Mathematical Geosciences*, February 2018, 50 (2), 209–234.

# Appendix C

## Supplementary Materials for *Innovation Bandits:*

### *A Dynamic Portfolio Strategy with Endogenous*

### *Rewards*

#### C.1 Model assumptions

The table below gives an overview of model assumptions made throughout this paper.

Name	Description
Diffusion: Independent processes	We assume that portfolio choices evolve as independent Itô diffusion processes.
Diffusion: Parameters	The time-inhomogeneous drift $\mu(t, X_t)$ and time-homogeneous diffusion coefficient $\sigma(X_t)$ are known for each diffusion process. They are locally bounded and smooth, such that each portfolio choice evolves as a continuous semimartingale. For the portfolio choice to be an innovation choice, the drift $\mu$ is also monotone and either strictly positive or strictly negative.
Diffusion: Local times	When we do not invest in a portfolio choice, the respective utility remains constant.
Utility: Properties	The utility function is smooth and decreasing. It is further bounded from above and

$$\mathbb{E}_{t,x} \left[ \int_t^\infty e^{-\alpha s} |u(X_i(s))| ds \right] < \infty$$

for all  $t \geq 0$  and technologies  $i$ . This assumption requires either the utility function  $u$  to be bounded or the exponential discount rate  $\alpha$  to be sufficiently large.

Utility: Time-linearity	We assume that utility is linear in time, subject to exponentially discounted time-preference.
Utility: Choice-linearity	We assume that utility is linear in our choices. The utility is a linear combination of the increasing utility associated with our portfolio choices.

Strategy: Investment feedback	We assume perfect historical information at any point in continuous time. For every investment made, we receive the outcome instantaneously and can update our strategy immediately.
Demand: Constant	We assume constant demand $D$ in each year (or general time interval). The cumulative deployment of a technology grows linearly while we invest in it.

Table C.1: Model assumptions made throughout this paper. The practical implications of these assumptions are described in the main text.

## C.2 Proofs: Gittins indices

In this appendix, we consider a broad class of diffusion processes which includes innovation bandits as defined in (V.5) and mean-reverting bandits as defined in (V.18). Without loss of generality, we assume the demand  $D = 1$  in the case of innovation bandits. Following (Milazzo, 2023), let  $X = (X(t))_{t \geq 0}$  be the time-inhomogeneous stochastic process solving

$$dX(t) = \mu(t, X(t))dt + \sigma(X(t))dB(t) \quad (\text{C.1})$$

in Itô's sense (Ito, 1951), or more precisely, satisfying

$$X(t) = X(0) + \int_0^t \mu(s, X(s))ds + \int_0^t \sigma(X(s))dB(s) \quad (\text{C.2})$$

for all  $t \geq 0$ , where the measurable functions  $\mu : [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$  and  $\sigma_i > 0$  are the drift and diffusion coefficients. Here,  $B$  is a standard Brownian motion started from 0. We always assume (2.2)–(2.4) of (Milazzo, 2024).

The utility function  $u : \mathbb{R} \rightarrow \mathbb{R}$  fulfills the assumptions of Table C.1, which includes the case of negative net present costs  $u(x) = -e^x$  from the main text.

As mentioned in Section V.3, we compute the Gittins index based on the retirement formulation. To this end, we recall the function  $L : \mathbb{R}^2 \rightarrow \mathbb{R}$  defined by

$$L(X(t), m) := \sup_{\tau \geq 0} \mathbb{E}_{t, X(t)} \left[ \int_0^\tau e^{-\alpha s} u(X(t+s))ds + me^{-\alpha \tau} \right], \quad (\text{C.3})$$

and the Gittins index defined by

$$M(t) = \inf\{m \in \mathbb{R} : L(X(t), m) = m\}. \quad (\text{C.4})$$

In words, the Gittins index is the smallest retirement reward  $m^*$  for which we are willing to stop immediately. We will thus be interested in studying the function  $L(x, m)$  and determining the optimal stopping time  $\tau^*$ . We refer the reader to [EL Karoui and Karatzas \(1994\)](#) for more background on these definitions.

The main ideas in what follows is to exploit the well-known connection between optimal stopping and free-boundary problems. For mean-reverting bandits, the computation of the Gittins index based on [Karatzas \(1984\)](#) is relatively straightforward since the Ornstein-Uhlenbeck process has time-homogeneous drift and diffusion coefficients. For innovation bandits, in which case it is essential to consider diffusion processes with time-inhomogeneous drift, the computation of the Gittins index is significantly more involved, and we will need to make use of recent advances in stochastic control theory and optimal stopping (e.g., [Peskir \(2005\)](#); [Peskir and Shiryaev \(2006\)](#); [de Angelis and Peskir \(2020\)](#), [de Angelis \(2015\)](#); [Milazzo \(2023\)](#)).

### C.2.1 Mean-reverting bandits

We follow Section 3 in [Karatzas \(1984\)](#). Our goal is to find  $\varphi(x)$  satisfying the following:

$$\frac{1}{2}\sigma^2\varphi''(x) + \theta(\kappa - x)\varphi'(x) - \alpha\varphi(x) = e^x \quad \text{on } (-\infty, b), \quad (\text{C.5})$$

$$\varphi(x) = m \quad \text{on } [b, \infty), \quad (\text{C.6})$$

$$\varphi'(b) = 0. \quad (\text{C.7})$$

First, a particular solution is given by

$$p(x) = -\mathbb{E}_x\left[\int_0^\infty e^{-\alpha t} e^{X(t)} dt\right]. \quad (\text{C.8})$$

Making use of the explicit representation of the Ornstein-Uhlenbeck process given by

$$X(t) = X(0)e^{-\theta t} + \kappa(1 - e^{-\theta t}) + \sigma \int_0^t e^{-\theta(t-s)} dB(s), \quad (\text{C.9})$$

one can alternatively rewrite Equation (C.8) as

$$p(x) = - \int_0^\infty e^{-\alpha t} e^{xe^{-\theta t} + \kappa(1 - e^{-\theta t})} \cdot \mathbb{E}_0 \left[ e^{\sigma \int_0^t e^{-\theta(t-s)} dB(s)} \right] dt, \quad (\text{C.10})$$

which facilitates the computation of  $p'(x)$ . Second, we solve the homogeneous equation

$$\frac{1}{2}\sigma^2 \phi''(x) + \theta(\kappa - x)\phi'(x) - \alpha\phi(x) = 0. \quad (\text{C.11})$$

To this end, we consider the ordinary Hermite equation

$$h''(x) - 2xh'(x) + 2vh(x) = 0, \quad (\text{C.12})$$

to which the couple  $H_v(\pm x)$  forms a fundamental solution<sup>1</sup>. Setting  $v := -\alpha/\theta$  and applying the change-of-variable  $y = \frac{\sigma x}{\sqrt{\theta}} + \kappa \iff x = \frac{\sqrt{\theta}}{\sigma}(y - \kappa)$  to Equation (C.12), we obtain

$$\frac{\sigma^2}{\theta} h'' \left( \frac{\sqrt{\theta}}{\sigma}(y - \kappa) \right) - 2(y - \kappa) h' \left( \frac{\sqrt{\theta}}{\sigma}(y - \kappa) \right) - 2\frac{\alpha}{\theta} h \left( \frac{\sqrt{\theta}}{\sigma}(y - \kappa) \right) = 0, \quad (\text{C.13})$$

$$\iff \frac{1}{2}\sigma^2 h'' \left( \frac{\sqrt{\theta}}{\sigma}(y - \kappa) \right) + \theta(\kappa - y) h' \left( \frac{\sqrt{\theta}}{\sigma}(y - \kappa) \right) - \alpha h \left( \frac{\sqrt{\theta}}{\sigma}(y - \kappa) \right) = 0. \quad (\text{C.14})$$

We conclude that the couple

$$H_{-\alpha/\theta} \left( \pm \frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right) \quad (\text{C.15})$$

forms a fundamental solution to Equation (C.11).

---

<sup>1</sup>We refer to Section 7 in [Alili et al. \(2005\)](#) for the definition of the Hermite function  $H_v$  and the relevant recurrence relations.

In summary, the general solution of Equation (C.5) is given by

$$\varphi(x) = A \cdot H_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right) + B \cdot H_{-\alpha/\theta} \left( -\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right) + p(x). \quad (\text{C.16})$$

Boundedness of  $\varphi$  on  $(-\infty, b)$  implies  $B = 0$ . From this, we conclude that due to (C.7)

$$M(x) = -\frac{p'(x)H_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right)}{\frac{d}{dx} \left( H_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right) \right)} + p(x) \quad (\text{C.17})$$

$$= -\frac{\sigma\sqrt{\theta}}{2\alpha} \frac{H_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right)}{H_{-\alpha/\theta-1} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right)} p'(x) + p(x), \quad (\text{C.18})$$

where we made use of the recurrence relation

$$\frac{\partial H_v}{\partial z}(z) = 2vH_{v-1}(z), \quad (\text{C.19})$$

which implies that

$$\frac{d}{dx} \left( H_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right) \right) = -\frac{\sqrt{\theta}}{\sigma} H'_{-\alpha/\theta} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right) \quad (\text{C.20})$$

$$= \frac{2\alpha}{\sigma\sqrt{\theta}} H_{-\alpha/\theta-1} \left( +\frac{\sqrt{\theta}}{\sigma}(\kappa - x) \right). \quad (\text{C.21})$$

## C.2.2 Innovation bandits

As mentioned above, the main ideas in what follows is to exploit the well-known connection between optimal stopping and free-boundary problems. We refer the reader to [Peskir and Shiryaev \(2006\)](#) for a comprehensive account of the general theory.

## A free-boundary problem

Firstly, we define a value function

$$\varphi(t, x) := \sup_{\tau \geq t} \mathbb{E}_{t,x} \left[ \int_t^\tau e^{-\alpha s} u(X(s)) ds + me^{-\alpha \tau} \right]. \quad (\text{C.22})$$

This represents the maximum utility in the interval  $[t, \infty)$  when starting at  $X(t) = x$ .

For every point  $(t, x)$ , we can now ask whether to continue or retire. This provides us with the disjoint *continuation set*  $C$  and *stopping set*  $D$ , where our respective stopping time is defined as

$$\tau_D := \inf\{t \geq 0 : X(t) \in D\}.$$

It is clear that for  $(t, x) \in C$ ,  $\varphi(t, x) > me^{-\alpha t}$ , and for  $(t, x) \in D$ ,  $\varphi(t, x) = me^{-\alpha t}$ .

Our task of finding the optimal stopping time  $\tau^*$  then reduces to finding the stopping boundary  $\partial D$  for a given  $m \in \mathbb{R}$ . A prominent way to solve this “free boundary problem” is to solve for  $\varphi(t, x)$ , which fulfills the following PDE (Øksendal, 2013, Chapter 10):

$$\begin{aligned} \frac{1}{2}\sigma(x)^2 \frac{\partial^2 \varphi}{\partial x^2} + \mu(t, x) \frac{\partial \varphi}{\partial x} + \frac{\partial \varphi}{\partial t} + u(x)e^{-\alpha t} = 0 \end{aligned} \quad \text{on } C, \quad (\text{C.23})$$

$$\left( \varphi > me^{-\alpha t} \quad \text{on } C, \right) \quad (\text{C.24})$$

$$\varphi = me^{-\alpha t} \quad \text{on } D, \quad (\text{C.25})$$

$$\varphi = 0 \quad \text{on } \{\infty\} \times \mathbb{R}, \quad (\text{C.26})$$

$$\frac{\partial \varphi}{\partial x} = 0 \quad \text{on } \partial D. \quad (\text{C.27})$$

We can use the results from Milazzo (2023) to simplify our search for the stopping boundary  $\partial D$ . Based on the differentiability and monotonicity of our drift parameters  $\mu(t, X)$  in (C.1), we find

**Lemma 4.** *For a portfolio choice following the assumptions in Table C.1, and retirement reward  $m \in \mathbb{R}$ , there exists a continuous and monotone function  $b_m : [0, \infty) \rightarrow \mathbb{R}$ , such*

that

$$C \equiv \{(t, x) \in [0, \infty) \times \mathbb{R} \mid x > b_m(t)\} \quad (\text{C.28})$$

$$D \equiv \{(t, x) \in [0, \infty) \times \mathbb{R} \mid x \leq b_m(t)\}. \quad (\text{C.29})$$

**Proof. Part I: Existence of the free boundary function**

For this proof, we first show the existence of  $b_m(t)$ . We follow the approach of [Jacka \(1991\)](#) to show that for  $(t, x) \in C$  (the continuation region) and  $y > x$ , then  $(t, y) \in C$ .

We first note that for  $(t, x) \in C$ ,  $\varphi(t, x) > me^{-\alpha t}$ . Moreover, for  $(t, x') \in D$ ,  $\varphi(t, x') = me^{-\alpha t}$ . We now fix  $(t, x) \in C$  with stopping time  $\tau = \inf\{s \geq t, X(s) = x : (X(s), s) \in D\}$  and consider  $y > x$ . Then

$$\varphi(y, t) - \varphi(x, t) = \varphi(y, t) - \mathbb{E}_{t,x} \left[ \int_t^\tau e^{-\alpha s} u(X(s)) ds + me^{-\alpha \tau} \right] \quad (\text{C.30})$$

$$\geq \mathbb{E}_{t,y} \left[ \int_t^\tau e^{-\alpha s} u(X(s)) ds + me^{-\alpha \tau} \right] - \mathbb{E}_{t,x} \left[ \int_t^\tau e^{-\alpha s} u(X(s)) ds + me^{-\alpha \tau} \right] \quad (\text{C.31})$$

$$= \mathbb{E}_t \left[ \int_t^\tau e^{-\alpha s} (u(X'(s)) - u(X(s))) ds \right], \quad (\text{C.32})$$

where  $X'(t) = y$  and  $X(t) = x$ .

Due to the monotonicity of  $u(X)$ , the integrand is strictly positive as long as  $X'(s) > X(s)$ . For  $X'(s) = X(s)$ , the expected value implies that the expression is zero. This means that

$$\varphi(y, t) \geq \varphi(x, t) \quad (\text{C.33})$$

$$> me^{-\alpha t}. \quad (\text{C.34})$$

**Part II: Simplifying the PDE**

Next, we consider the monotonicity of  $b_m(t)$ . For this, we apply the result of [Milazzo \(2023\)](#), but we first need to re-formulate our PDE (C.23)-(C.27) as an optimal stopping problem without path dependence.

It is convenient to define  $h : [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$  as

$$h(t, x) := \mathbb{E}_{t,x} \left[ \int_t^\infty e^{-\alpha s} u(X(s)) ds \right], \quad (\text{C.35})$$

which satisfies  $\|h\|_\infty \leq e^{-\alpha t} \frac{\|u\|_\infty}{\alpha}$ .  $h$  is the reward corresponding to the stopping time  $\tau = \infty$  (a.s.). Importantly,  $h$  is a particular solution to the PDE (C.23).

Using  $h$ , let's define a new value function

$$v(t, x) := e^{-\alpha t} \varphi(t, x) - h(t, x) = \sup_{\tau \geq t} \mathbb{E}_{t,x} \left[ - \int_\tau^\infty e^{-\alpha s} u(X(s)) ds + m e^{-\alpha \tau} \right] \quad (\text{C.36})$$

$$= \sup_{\tau \geq t} \mathbb{E}_{t,x} \left[ - \mathbb{E}_{t,x} \left[ \int_\tau^\infty e^{-\alpha s} u(X(s)) ds \mid \tau, X_\tau \right] + m e^{-\alpha \tau} \right] \quad (\text{C.37})$$

$$= \sup_{\tau \geq t} \mathbb{E}_{t,x} \left[ - \mathbb{E}_{\tau, X_\tau} \left[ \int_\tau^\infty e^{-\alpha s} u(X(s)) ds \right] + m e^{-\alpha \tau} \right] \quad (\text{C.38})$$

$$= \sup_{\tau \geq t} \mathbb{E}_{t,x} \left[ -h(\tau, X_\tau) + m e^{-\alpha \tau} \right] \quad (\text{C.39})$$

$$= \sup_{\tau \geq t} \mathbb{E}_{t,x} [G(\tau, X_\tau)], \quad (\text{C.40})$$

where  $G(s, y) := -h(s, y) + m e^{-\alpha s} = -\mathbb{E}_{s,y} \left[ \int_s^\infty e^{-\alpha u} u(X(u)) du \right] + m e^{-\alpha s}$ , and we used the strong Markov property to remove the integral process.

We can see that  $(\varphi, b)$  now solves the free-boundary problem (Equations C.23-C.27) if and only if  $(v, b_m)$  solves the free-boundary problem

$$\frac{1}{2} \sigma(x)^2 \frac{\partial^2 v}{\partial x^2} + \mu(t, x) \frac{\partial v}{\partial x} + \frac{\partial v}{\partial t} = 0 \quad \text{on } C, \quad (\text{C.41})$$

$$\left( v > G \quad \text{on } C, \right) \quad (\text{C.42})$$

$$v = G \quad \text{on } D, \quad (\text{C.43})$$

$$v = 0 \quad \text{on } \{\infty\} \times \mathbb{R}, \quad (\text{C.44})$$

$$\frac{\partial v}{\partial x} = \frac{\partial G}{\partial x} \quad \text{on } \partial D. \quad (\text{C.45})$$

**Remark 5.** *The value of  $v$  at  $[0, \infty) \times \{\infty\}$  equals  $e^{-\alpha t} \max\{0, m - \sup_{x \in \mathbb{R}} u(x)\}$ . This boundary condition could also be added to the problem.*

### Part III: Monotonicity of the free boundary

In order to show the monotonicity of  $b_m(t)$ , we need to show that Assumption 2.1 of [Milazzo \(2023\)](#) holds. Specifically, we need to show that  $G(s, y)$  is upper semi-continuous and that

$$\mathbb{E}_{t,x} \left[ \sup_{s \geq t} |G(s, X_s)| \right] < \infty.$$

Upper semi-continuity is trivial for our case of  $G(s, y)$  due to the continuity in  $X$  and  $u$ .

The boundedness of  $G$  will depend on the boundedness of  $h(s, y)$  since  $\alpha > 0$ . For bounded utility  $u(X) > 0$ , this is given trivially. We can relax this condition by assuming  $h(t, X_t) < \infty$  instead. This also allows for unbounded utility functions as long as  $\alpha$  is sufficiently large.

### Part IV: Continuity of the free boundary

To show the continuity of  $b_m(t)$ , we distinguish discontinuities of the first and second kinds.

For the second kind, we note that  $b_m(t)$  is finite due to the continuity and boundedness of  $\mu(t, X)$ ,  $\sigma(X)$ , and  $G(t, X)$ . Since  $b_m(t)$  is also monotone, there are no discontinuities of the second kind. For the first kind, we use Theorem 3 of [Peskir \(2019\)](#). We can write our time inhomogeneous stochastic process  $X_t$  as a two-dimensional (homogeneous) space-time diffusion process. The infinitesimal generator of this process is elliptic, so we can apply Theorem 3 of [Peskir \(2019\)](#). Due to the smooth-fit condition in Equation (C.45), this implies that  $b_m(t)$  has no discontinuities of the first kind. Therefore,  $b_m(t)$  is continuous.  $\square$

A common way to find  $b_m(t)$  is to solve the PDE (C.23)-(C.27), such as in [Karatzas \(1984\)](#). However, this is generally difficult and usually requires stronger conditions on  $\sigma$  and  $\mu$  than those we have here. Fortunately, as we will see in the next section, we can still derive an implicit equation for  $b_m(t)$  that is sufficient to calculate the Gittins index without finding the full solution to the free boundary problem.

## Integral equation for the Gittins index

The key idea to find an expression for  $b(t) = b_m(t)$  is the use of the change-of-variable formula from [Peskir \(2005\)](#). This formula is a generalization of the Itô-Tanaka formula applied to  $\varphi(x, t)$  after some algebraic operations. We can then apply our boundary conditions to find the free-boundary equation for  $b(t)$  in [Theorem 1](#).

**Remark 6.** *As we can see, this expression now depends only on the free boundary  $b(t)$  and the stochastic process  $X(t)$  itself. Moreover, for any time  $t$ , the boundary depends exclusively on future values of the utility and boundary at  $t + u$ .*

*The expression in the expected value is equivalent to the (marginal) additional utility of continuing the process relative to retiring.  $m\alpha$  is the marginal utility of retiring the process, with the  $\alpha$  factor coming from the exponential discounting. The indicator function  $\mathbf{1}_{X(t+u) > b(t+u)}$  conditions the expected value on the continuation region  $C$ . In other words, when on the boundary  $\partial D$ , the expected future utility within the continuation region equals the retirement utility outside the continuation region.*

### *Proof Theorem 1. Part I: Change of variables*

In this first step, we apply the change-of-variable formula from [Peskir \(2005\)](#) to the PDE [\(C.41\)](#)-[\(C.45\)](#). The change-of-variable formula is a generalization of the Itô-Tanaka formula.

Assume that  $b : [0, \infty) \rightarrow \mathbb{R}$  and  $v : [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$  are continuous and satisfy the additional assumptions:

1.  $v$  is  $C^{1,2}$  on  $C$  and on  $D$ .
2.  $b$  continuous and of bounded variation.
3.  $\frac{1}{2}\sigma(t, x)^2 \frac{\partial^2 G}{\partial x^2} + \mu(t, x) \frac{\partial G}{\partial x} + \frac{\partial G}{\partial t}$  locally bounded on  $D$ .
4.  $\frac{\partial v}{\partial x}$  is continuous on  $\bar{C}$  and  $\bar{D}$ .
5.  $v$  and  $b$  are monotone.

Then, we have

$$\begin{aligned}
v(t, X(t)) &= v(0, X(0)) + \int_0^t \left( \frac{1}{2} \sigma^2 \frac{\partial^2 v}{\partial x^2} + \mu \frac{\partial v}{\partial x} + \frac{\partial v}{\partial t} \right) (s, X(s)) \mathbf{1}_{X(s) \neq b(s)} ds \\
&+ \int_0^t \left( \sigma(x) \frac{\partial v}{\partial x} \right) (s, X(s)) \mathbf{1}_{X(s) \neq b(s)} dB(s) \\
&+ \frac{1}{2} \int_0^t \left( \frac{\partial v}{\partial x}(s, X(s)+) - \frac{\partial v}{\partial x}(s, X(s)-) \right) \mathbf{1}_{X(s)=b(s)} d\ell_X^b(s), \tag{C.46}
\end{aligned}$$

where  $\ell_X^b(s)$  is the local time of  $X$  at the curve  $b$ . Here, we write  $\frac{\partial v}{\partial x}(t, x+)$  for the right  $x$ -derivative at  $(t, x)$  and  $\frac{\partial v}{\partial x}(t, x-)$  for the left  $x$ -derivative at  $(t, x)$ .

In this formulation, we found the integral representation of our solution  $v$ . This can be translated into an integral formulation for  $b$  in the third step. We are, therefore, left to show assumptions 1 – 5.

Assumptions 1 and 4. The continuity and smoothness of  $v$  on  $\bar{C}$  follow from the continuity of  $X$  and  $G$  (see [Peskir and Shiryaev \(2006, Chapter III\)](#) and [de Angelis \(2015\)](#).)

Assumptions 2 and 5 follow from Lemma 4.

Assumption 3 simply follows from the fact that all elements in the sum are bounded.

## Part II: Free boundary equation

Following Chapter 14.1 in [Peskir and Shiryaev \(2006\)](#), we use the previous integral formulation for  $v$  to derive the free-boundary equation for  $b$ .

By the smooth-fit condition (Equation C.45), the last integral in Equation C.46 vanishes. By Equation C.41, the first integral can be simplified. Together, we obtain

$$\begin{aligned}
v(t+s, X(t+s)) &= v(t, x) \\
&+ \int_0^s \left( \frac{1}{2} \sigma^2 \frac{\partial^2 G}{\partial x^2} + \mu \frac{\partial G}{\partial x} + \frac{\partial G}{\partial t} \right) (t+u, X(t+u)) \mathbf{1}_{X(t+u) < b(t+u)} du \\
&+ \int_0^s \left( \sigma \frac{\partial v}{\partial x} \right) (t+u, X(t+u)) \mathbf{1}_{X(t+u) \neq b(t+u)} dB(t+u). \tag{C.47}
\end{aligned}$$

Taking expectations on both sides, we get

$$\begin{aligned} \mathbb{E}_{t,x} [v(t+s, X(t+s))] &= v(t, x) \\ &+ \int_0^s \mathbb{E}_{t,x} \left[ \left( \frac{1}{2} \sigma^2 \frac{\partial^2 G}{\partial x^2} + \mu \frac{\partial G}{\partial x} + \frac{\partial G}{\partial t} \right) (t+u, X(t+u)) \mathbf{1}_{X(t+u) < b(t+u)} \right] du, \end{aligned} \quad (\text{C.48})$$

where we used that the second integral is a martingale starting at 0.

Assuming that  $\mathbb{E}_{t,x} [v(t+s, X(t+s))] \rightarrow 0$  as  $s \rightarrow \infty$  and using that  $h$  is a particular solution to Equation C.23, we deduce that

$$\begin{aligned} v(t, x) &= \int_0^\infty m\alpha e^{-\alpha(t+u)} \cdot \mathbb{E}_{t,x} [X(t+u) < b(t+u) \mathbf{1}_{X(t+u) < b(t+u)}] du \\ &\quad - \int_0^\infty e^{-\alpha(t+u)} \cdot \mathbb{E}_{t,x} [u(X(t+u)) \mathbf{1}_{X(t+u) < b(t+u)}] du. \end{aligned} \quad (\text{C.49})$$

Finally, setting  $x = b(t)$  and using the boundary condition of Equation C.43, we get the following for every  $t \in [0, \infty)$ :

$$\begin{aligned} me^{-\alpha t} - h(t, b(t)) &= \int_0^\infty m\alpha e^{-\alpha(t+u)} \cdot \mathbb{E}_{t,b(t)} [X(t+u) < b(t+u) \mathbf{1}_{X(t+u) < b(t+u)}] du \\ &\quad - \int_0^\infty e^{-\alpha(t+u)} \cdot \mathbb{E}_{t,b(t)} [u(X(t+u)) \mathbf{1}_{X(t+u) < b(t+u)}] du. \end{aligned} \quad (\text{C.50})$$

Equivalently, for every  $t \in [0, \infty)$ , we obtain the free boundary equation.

□

### Translational invariance of the free boundary

In this subsection, we assume that the stochastic process  $(X(t))_{t \geq 0}$  is translational invariant, meaning that the law of  $(X(t+s) + a)_{s \geq 0}$  started from  $X(t) = x$  is equal to the law of  $(X(t+s))_{s \geq 0}$  started from  $X(t) = x + a$ . We note that the diffusion process defined in (C.1) is translation invariant if the drift and diffusion coefficients do not depend on the process, i.e.  $\mu : [0, \infty) \rightarrow \mathbb{R}$  and  $\sigma > 0$  is constant. This is the case for innovation bandits defined in (V.5).

The translational invariance of the free boundary  $b_m$  in Corollary 2 then follows directly from the integral representation in Equation (V.11).

*Proof Corollary 2.* We first note that for  $b_1(t)$ , we have

$$\int_0^\infty e^{-\alpha s} \mathbb{E}_{t, b_1(t)} [(u(X(t+s)) - \alpha) \mathbf{1}_{X(t+s) > b_1(t+s)}] ds = 0. \quad (\text{C.51})$$

Dividing Equation (V.11) by  $m$  and using the translational property of  $u(\cdot)$ , we get for  $b_m(t)$

$$\int_0^\infty e^{-\alpha s} \mathbb{E}_{t, b_m(t)} [(u(X(t+s) + f(m^{-1})) - \alpha) \mathbf{1}_{X(t+s) > b_m(t+s)}] ds = 0. \quad (\text{C.52})$$

Adding  $f(m^{-1})$  to both sides of the inequality of the indicator function, we obtain

$$\int_0^\infty e^{-\alpha s} \mathbb{E}_{t, b_m(t)} [(u(X(t+s) + f(m^{-1})) - \alpha) \mathbf{1}_{X(t+s) + f(m^{-1}) > b_m(t+s) + f(m^{-1})}] ds = 0. \quad (\text{C.53})$$

Finally, using that the law of  $(X(t+s) + f(m^{-1}))_{s \geq 0}$  started from  $X(t) = b_m(t)$  is equal to the law of  $(X(t+s))_{s \geq 0}$  started from  $X(t) = b_m(t) + f(m^{-1})$ , we get

$$\int_0^\infty e^{-\alpha s} \mathbb{E}_{t, b_m(t) + f(m^{-1})} [(u(X(t+s)) - \alpha) \mathbf{1}_{X(t+s) > b_m(t+s) + f(m^{-1})}] ds = 0. \quad (\text{C.54})$$

Since  $b_m(t)$  is unique for all  $m \in \mathbb{R}$ , we conclude that  $b_1(t) = b_m(t) + f(m^{-1})$ .  $\square$

## References

- Alili, L., P. Patie, and Jesper Lund Pedersen**, “Representations of the First Hitting Time Density of an Ornstein-Uhlenbeck Process 1,” *Stochastic Models*, October 2005, *21* (4), 967–980.
- de Angelis, Tiziano**, “A Note on the Continuity of Free-Boundaries in Finite-Horizon Optimal Stopping Problems for One-Dimensional Diffusions,” *SIAM Journal on Control and Optimization*, January 2015, *53* (1), 167–184.
- **and G. Peskir**, “Global  $C^1$  Regularity of the Value Function in Optimal Stopping Problems,” *The Annals of Applied Probability*, June 2020, *30* (3).
- Ito, Kiyoshi**, “III. Stochastic Differential Equation and Stochastic Integral Equation,” in “On Stochastic Differential Equations,” Vol. 1, United States: American Mathematical Society, 1951.
- Jacka, S. D.**, “Optimal Stopping and the American Put,” *Mathematical Finance*, 1991, *1* (2), 1–14.
- Karatzas, Ioannis**, “Gittins Indices in the Dynamic Allocation Problem for Diffusion Processes,” *The Annals of Probability*, February 1984, *12* (1), 173–192.
- Karoui, Nicole EL and Ioannis Karatzas**, “Dynamic Allocation Problems in Continuous Time,” *The Annals of Applied Probability*, May 1994, *4* (2), 255–286.
- Milazzo, Alessandro**, “On the Monotonicity of the Stopping Boundary for Time-Inhomogeneous Optimal Stopping Problems,” January 2023.
- , “On the Monotonicity of the Stopping Boundary for Time-Inhomogeneous Optimal Stopping Problems,” *Journal of Optimization Theory and Applications*, October 2024, *203* (1), 336–358.

**Øksendal, Bernt Karsten**, *Stochastic Differential Equations : An Introduction with Applications* Universitext, sixth edition, sixth corrected printing ed., Heidelberg: Springer, 2013.

**Peskir, Goran**, “A Change-of-Variable Formula with Local Time on Curves,” *Journal of Theoretical Probability*, July 2005, 18 (3), 499–535.

– , “Continuity of the Optimal Stopping Boundary for Two-Dimensional Diffusions,” *The Annals of Applied Probability*, February 2019, 29 (1).

– **and Albert Shiryaev**, *Optimal Stopping and Free-Boundary Problems* Lectures in Mathematics ETH Zürich, Basel ; Boston: Birkhäuser Verlag, 2006.