Supplementary Information for Probabilistic projections of global wind and solar power growth based on historical national experience

Avi Jakhmola[1], Jessica Jewell[1,2,3], Vadim Vinichenko[1], Aleh Cherp[4,5*]

[1*]Division of Physical Resource Theory, Department of Space, Earth and Environment, Chalmers University of Technology, Gothenburg, 41296, Sweden.
[2]Centre for Climate and Energy Transformations and Department of Geography, University of Bergen, Bergen, 5020, Norway.
[3]Advancing Systems Analysis, International Institute for Applied Systems Analysis, Laxenburg, 2361, Austria.
[4]Department of Environmental Science and Policy, Central European University University, Vienna, 1100, Austria.
[5]International Institute for Industrial Environmental Economics, Lund University, Lund, 22100, Sweden.

*Corresponding author(s). E-mail(s): cherpa@ceu.edu;

Contents

Supplementary Figures

Supplementary Figure 1 National and global takeoff thresholds.
Faint blue dots show takeoff thresholds for individual countries, the boxplots highlight their median, IQR and ranges; red dots indicate global takeoff thresholds. Note: the horizontal axis is on a logarithmic scale.

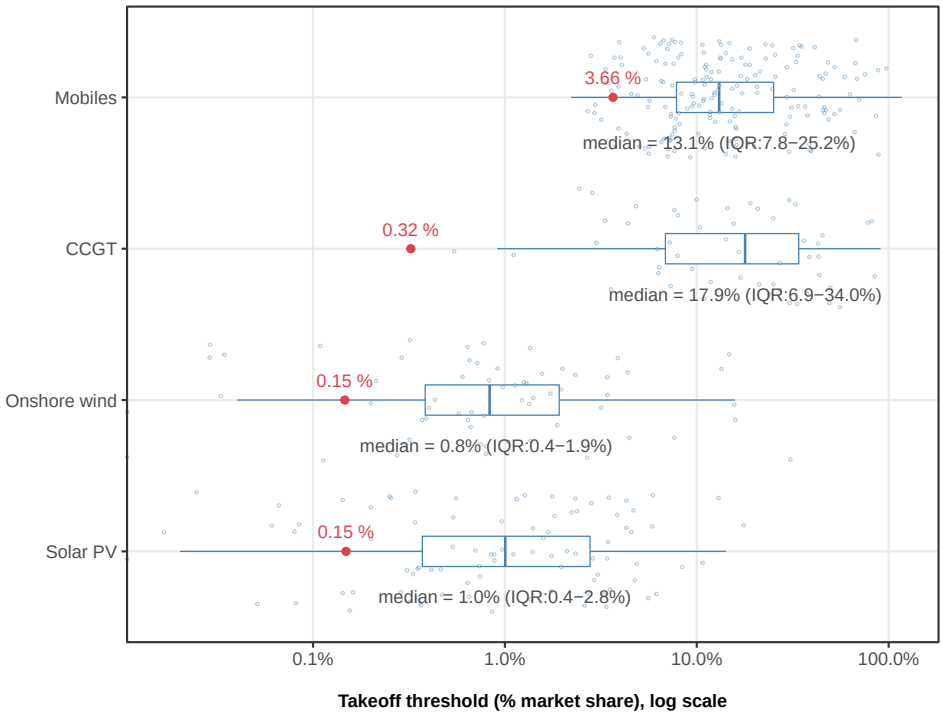

Supplementary Figure 2 Spatial diffusion of technologies.
The cumulative number of countries in different phases of technology growth at different years and the diffusion durations for different technologies. Bars show the number of countries with no deployment (gray), in the formative phases (light blue), in the acceleration phase (dark blue), and in or beyond the stable growth phase (yellow); the gray vertical line indicates the year of global takeoff.

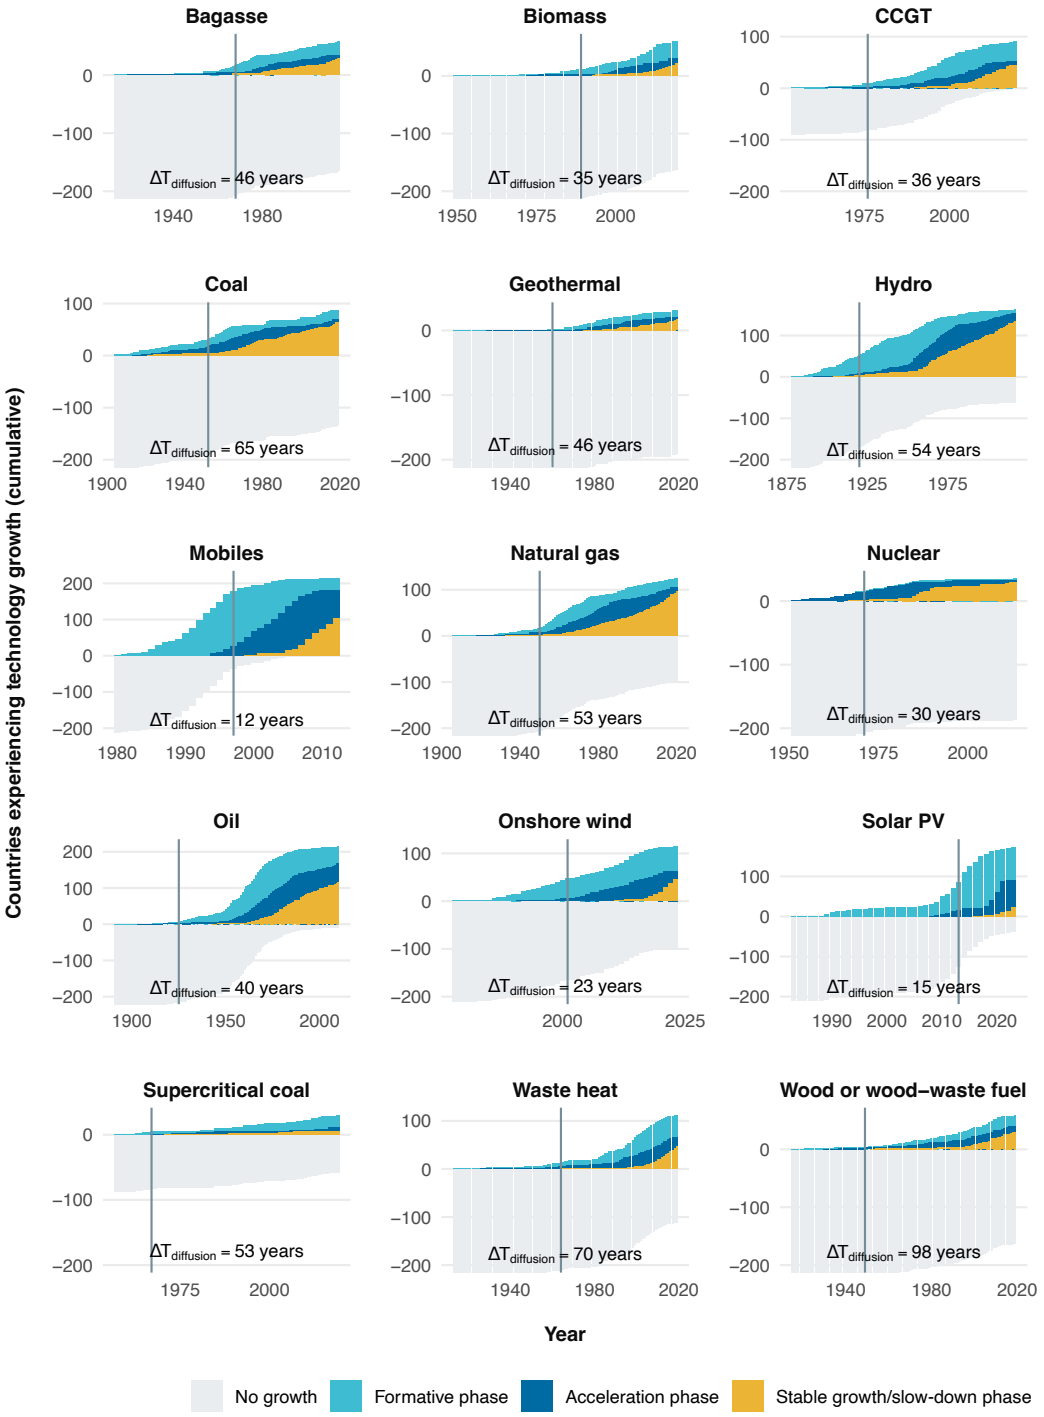

Supplementary Figure 3 Historical technology deployment curves.
Solid lines show global deployment measured in absolute units (green) and as percentage market share (blue); stars indicate takeoff.

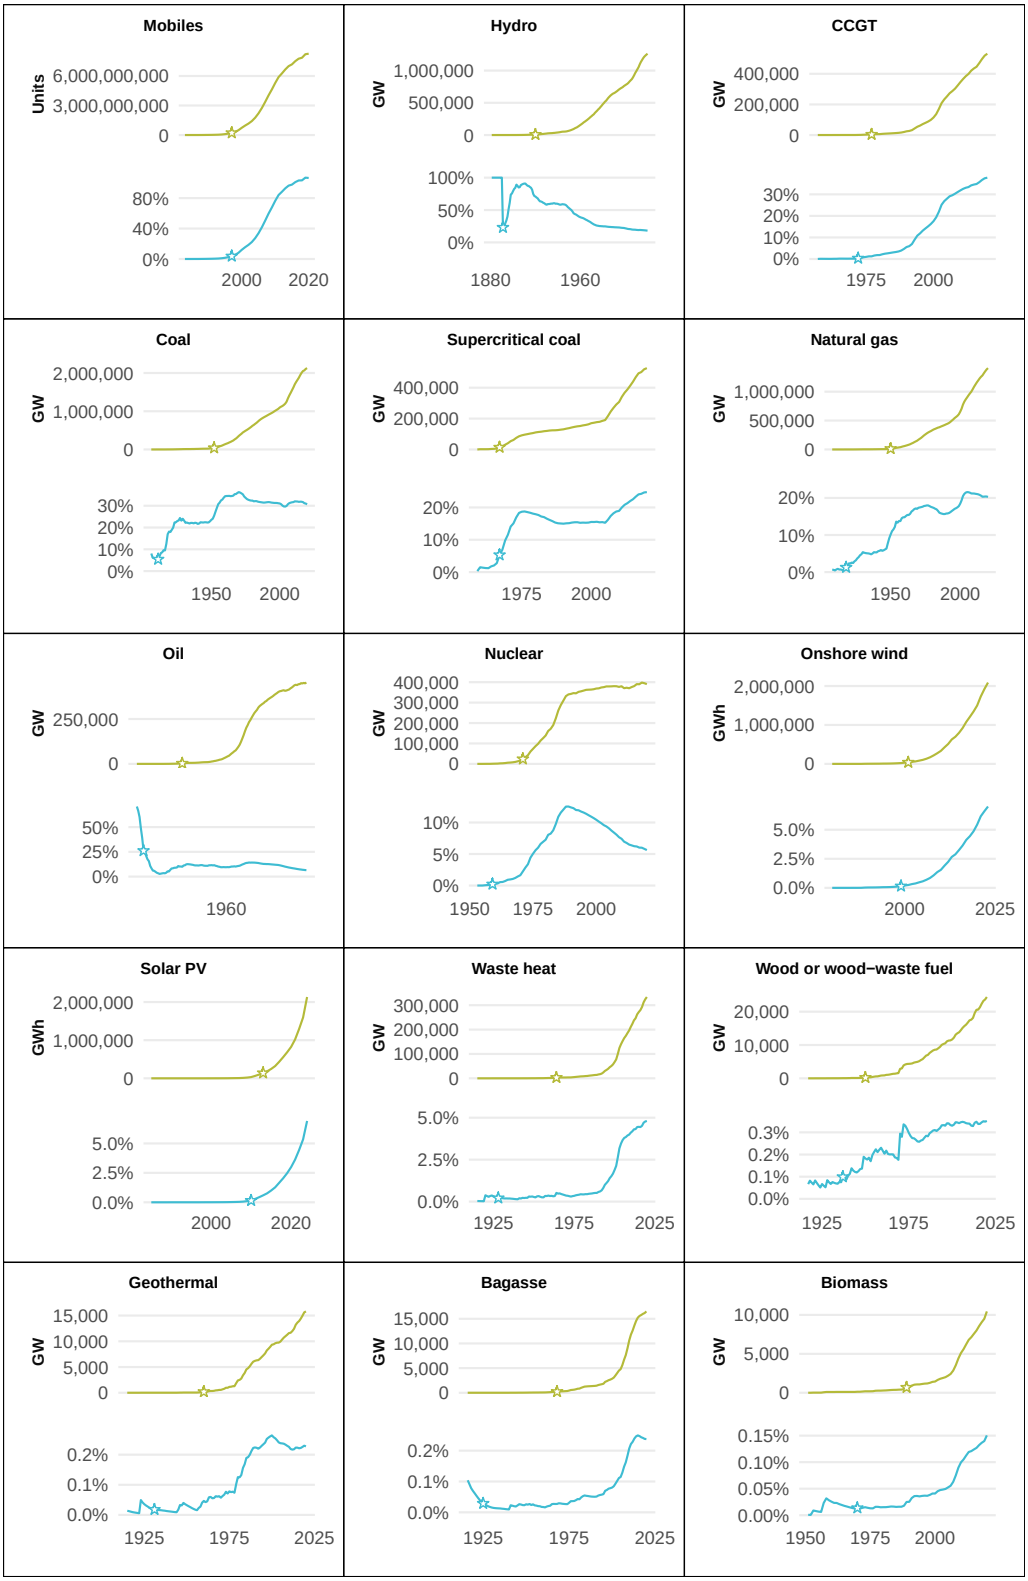

Supplementary Figure 4 Historical annual technology additions.
Bars show global annual additions measured in absolute units (green) and as percentage market share (blue) with Chinese shares highlighted in darker shades, solid lines show their three-year trailing average; stars indicate takeoff; red vertical dashed lines mark inflection points for logistic functions fit to the full deployment time-series.

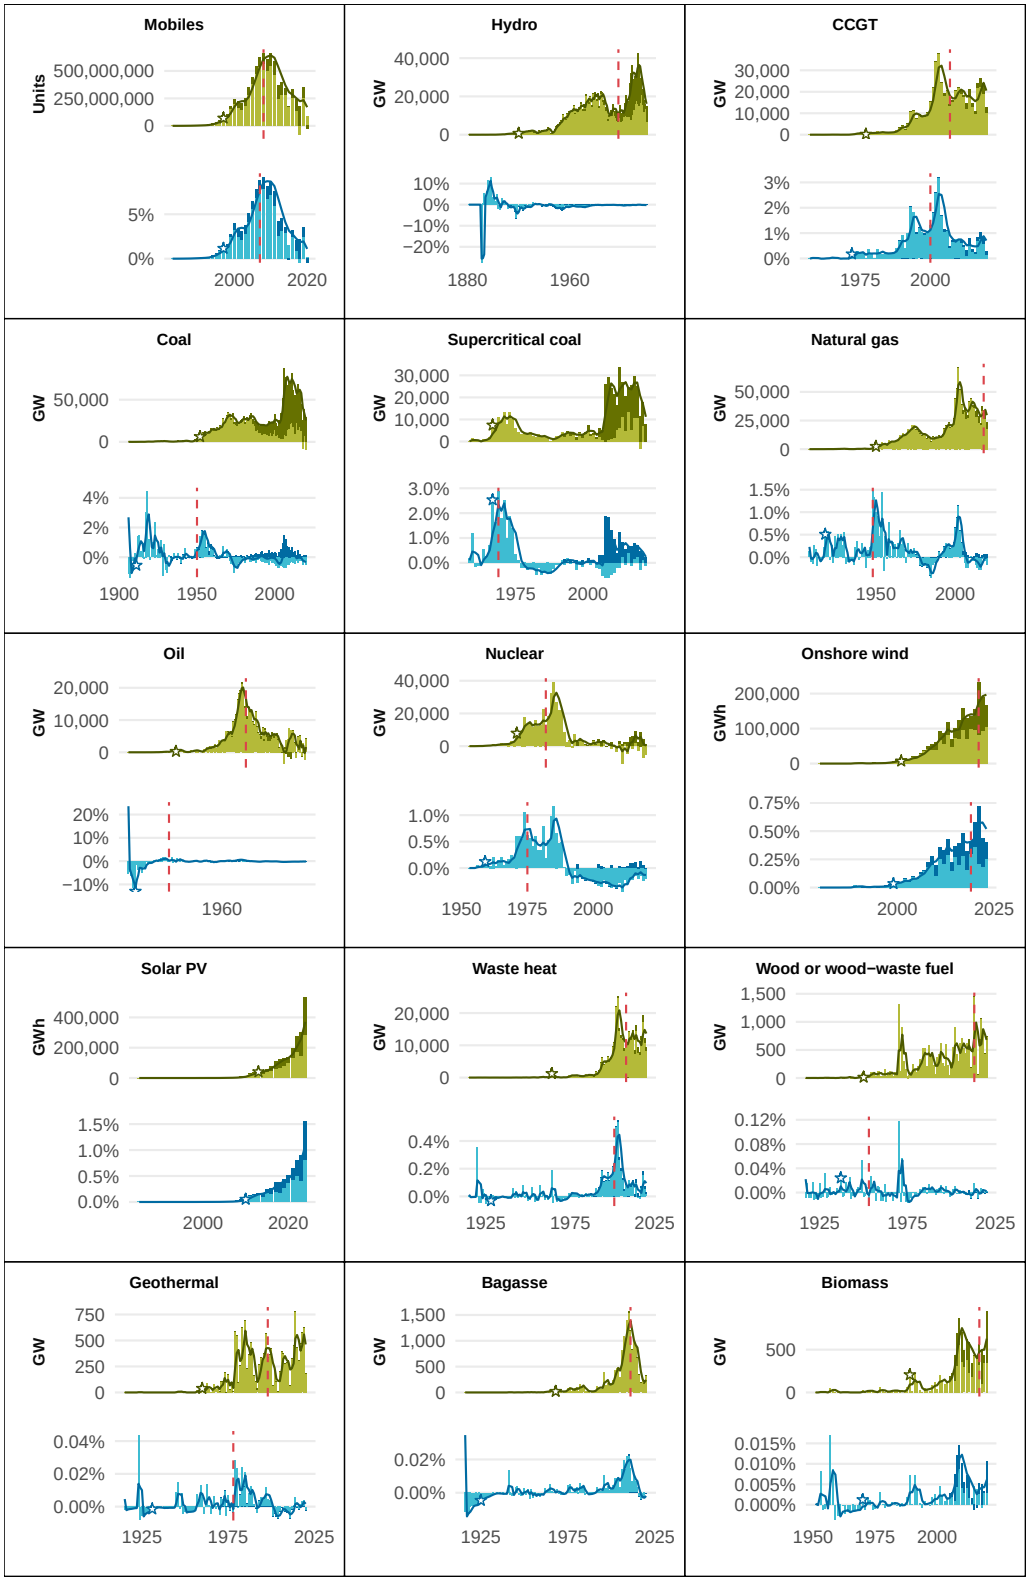

Supplementary Figure 5 The 'moving' L problem.
(a) Illustration of how logistic curves fit to incomplete deployment data struggle to correctly estimate the technology deployment ceiling, L. (b) Estimates for the deployment ceiling "L" (red squares) from logistic curves fit to empirical deployment data (solid gray line) curtailed at different years.

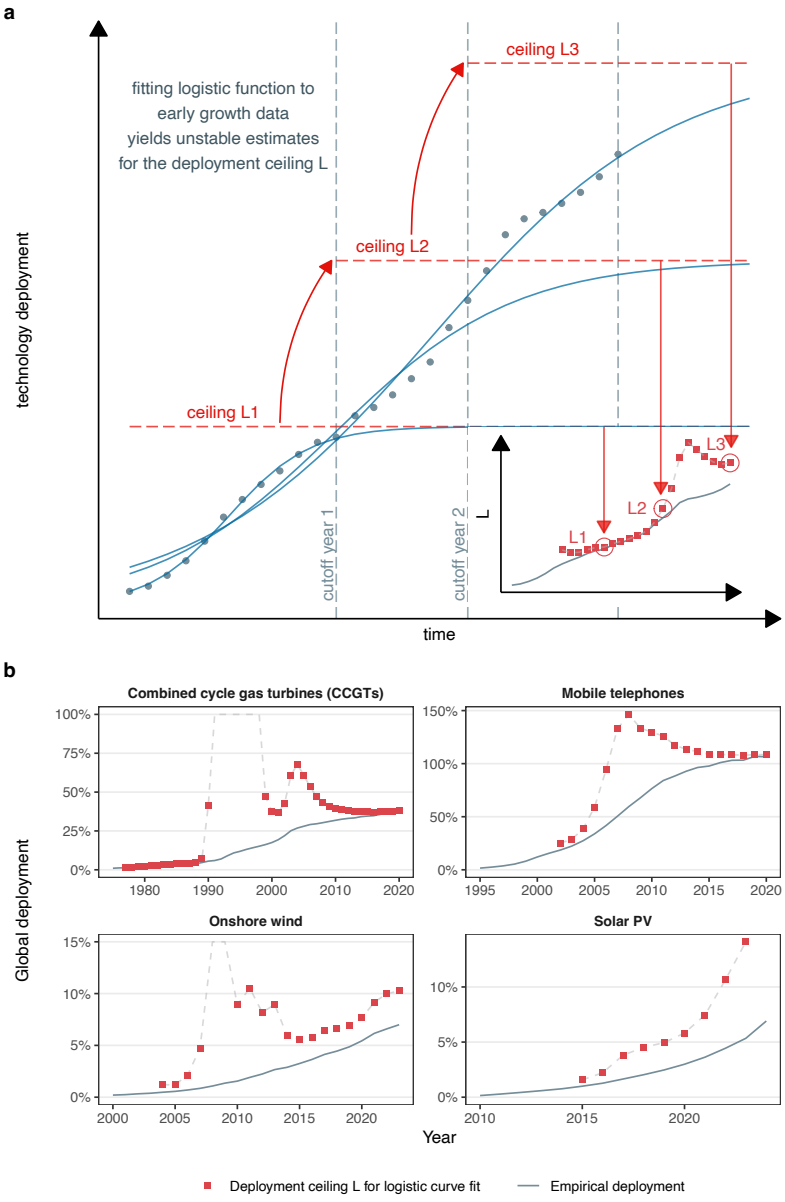

Supplementary Figure 6 The evolution of national and global logistic curve parameters.
Gray bars show historical values and gray dashed lines show their three-year trailing average. Red dots show parameters for logistic curves fit to global data. Solid blue lines and shaded blue intervals show the median and IQR for a sample of parameters for logistic curves fit to national data. Top row shows the global year-on-year growth rate, and the logistic intrinsic growth rate (k). Middle row shows the global deployment, and the logistic growth asymptote (L). Bottom row shows global annual additions, and the logistic maximum annual growth rate at the inflection point (G)

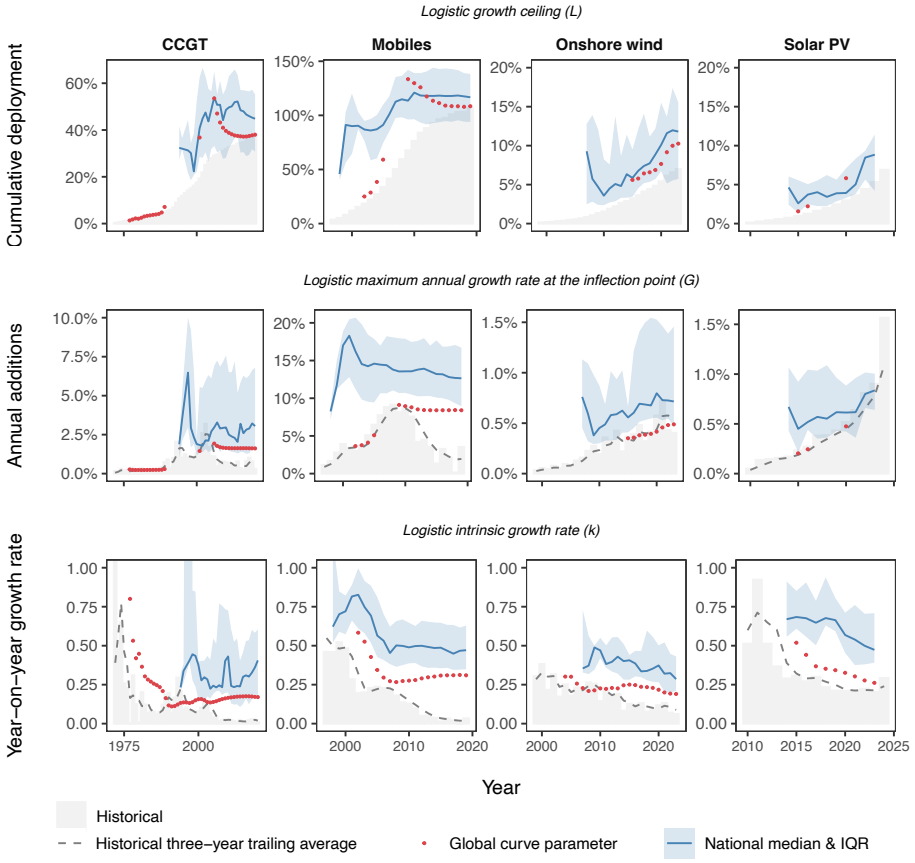

Supplementary Figure 7 Accuracy indicators for different variants of the PROLONG model. Indicators measuring the projection performance of the PROLONG model trained on simulations assuming individuals countries follow only bilogistic functions (yellow), only logistic functions (olive), or are equally likely to follow either function ('Mixed', green). The symmetric Mean Absolute Percentage Error (a) and symmetric Mean Percentage Error (b) for the median global projections from each model variant at different curtail years. The mean Continuous Ranked Probability score (c) for each model variant for different forecast horizons.

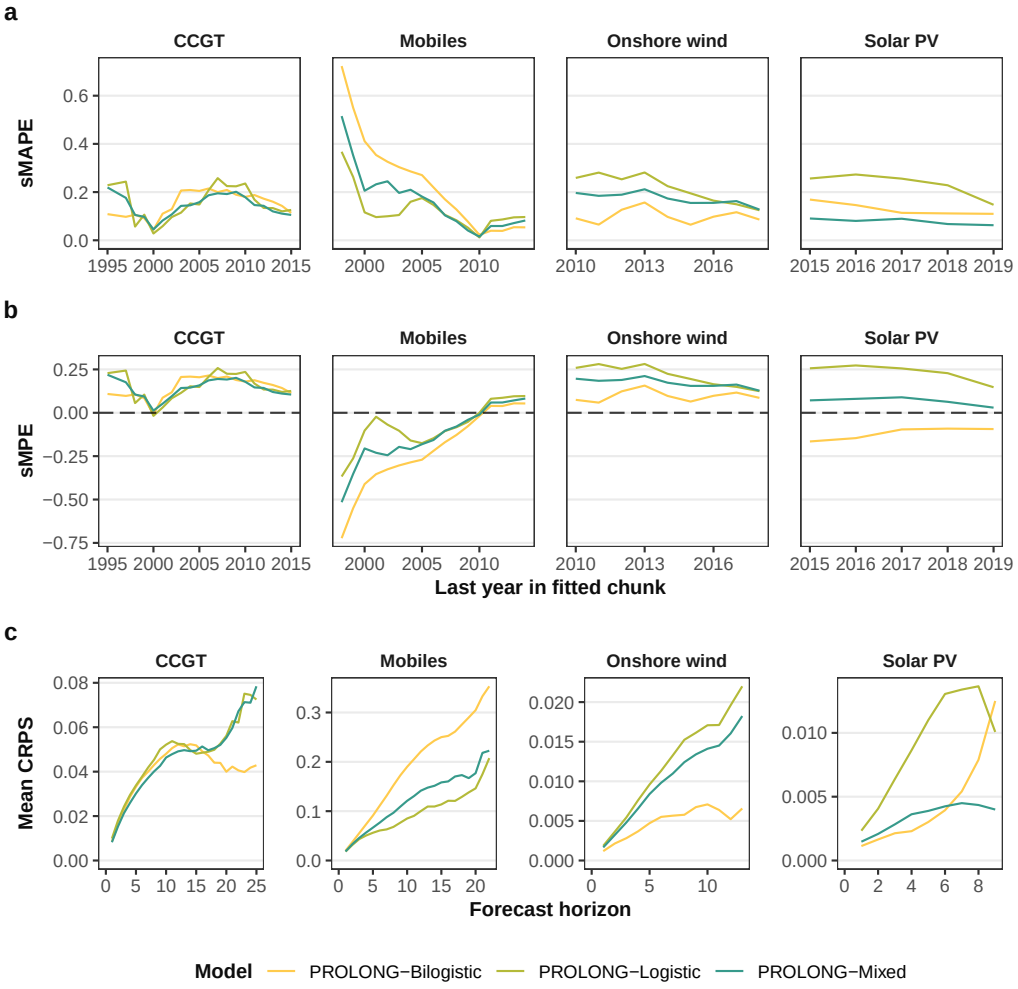

Supplementary Figure 8 Calibration of projection intervals for different variants of the PROLONG model

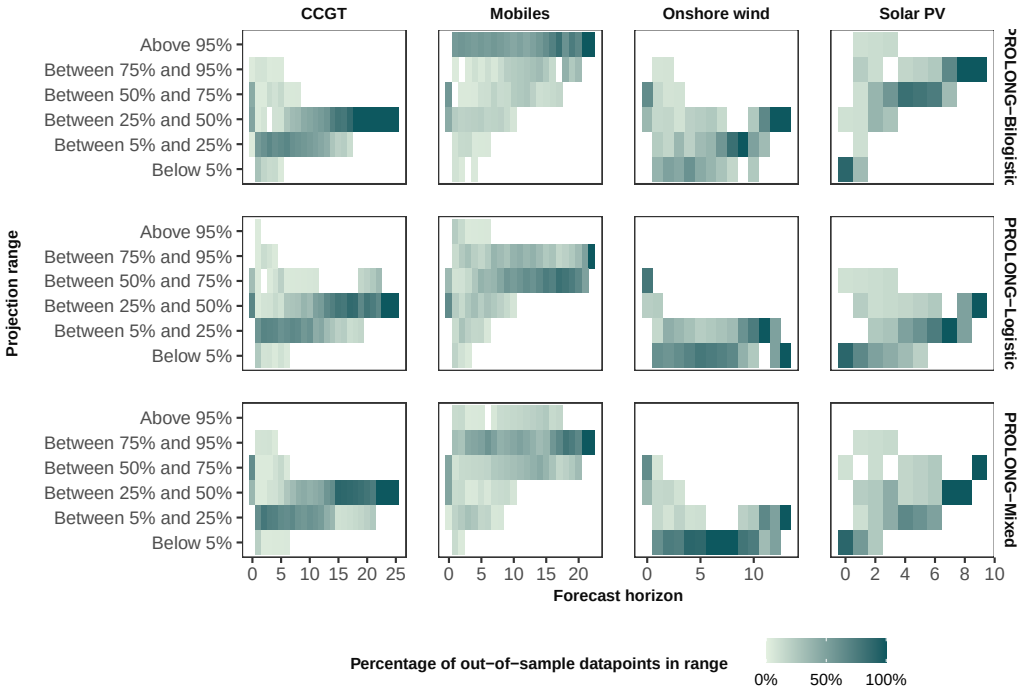

Supplementary Figure 9 Comparing the probabilistic projections for onshore wind and solar PV growth between 2023-2050 to current policy scenarios and market forecasts
Black dots show historical data; solid lines show the median, dark shaded intervals the 25-75th percentile, and light shaded intervals the 5-95th percentile projections from PROLONG; stars show deployment aligned with the COP28 Global Pledge on Renewables; crosses show forecasts from the IEA Renewables 2024 report; dot-dash lines show the IEA Stated Policies scenario (green), the BloombergNEF Economic Transition scenario (red), and three Current Policies scenarios from the IPCC AR6 report (magenta, blue, olive).

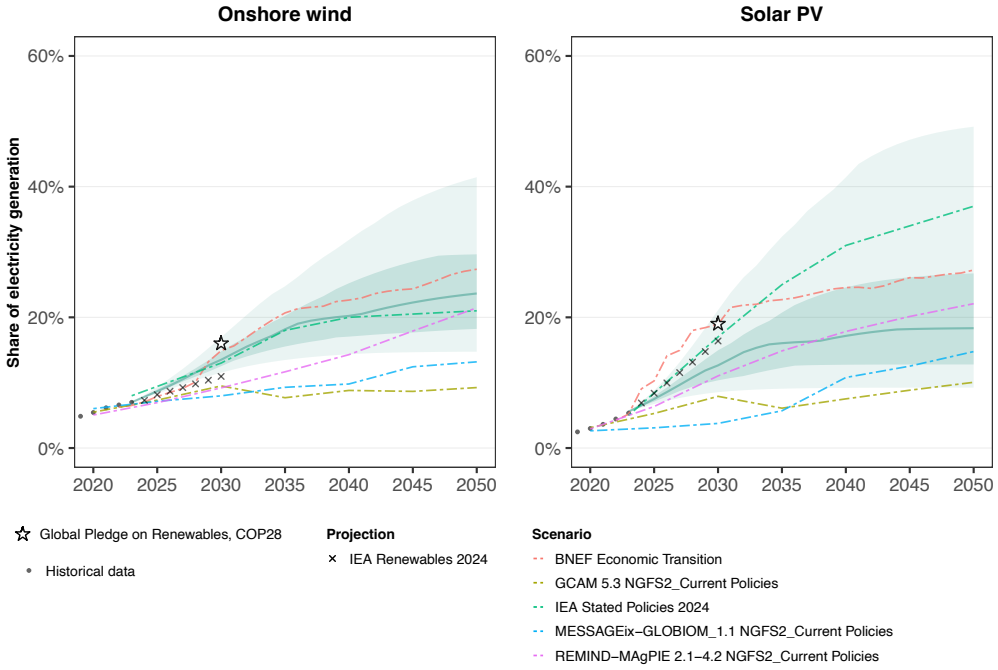

Supplementary Figure 10 Comparing the probabilistic projections for onshore wind and solar PV growth between 2023-2050 to scenarios from the literature.
Black dots show historical data; solid lines show the median, dark shaded intervals the 25-75th percentile, and light shaded intervals the 5-95th percentile projections from PROLONG; stars show deployment aligned with the COP28 Global Pledge on Renewables; dot-dash lines show deployment in refs. [1–5].

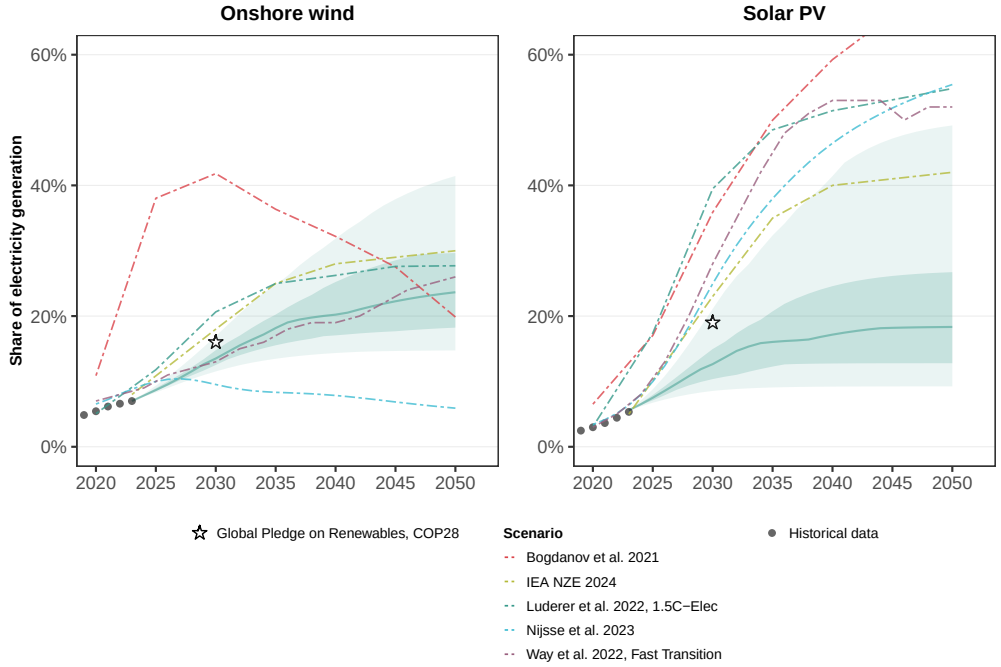

Supplementary Figure 11 Region-wise onshore wind deployment under the Baseline, Early and Late acceleration scenarios.
Dots show historical data; solid lines show deployment in the Basline (red), Early acceleration (green) and Late acceleration scenarios (blue).

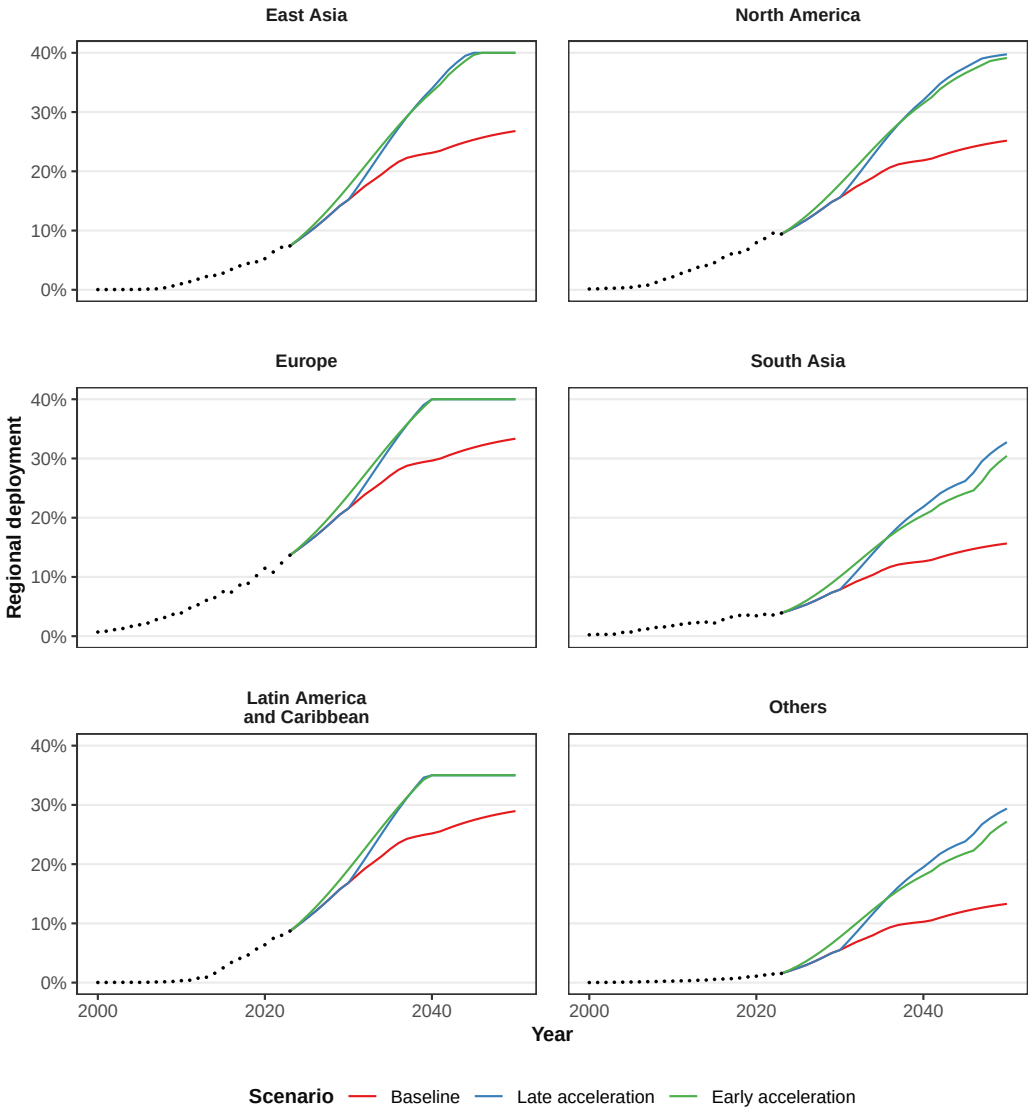

Supplementary Figure 12 Region-wise annual onshore wind additions under the Baseline, Early and Late acceleration scenarios.
Gray bars show historical data; coloured bars show deployment in the Basline (red), Early acceleration (green) and Late acceleration scenarios (blue).

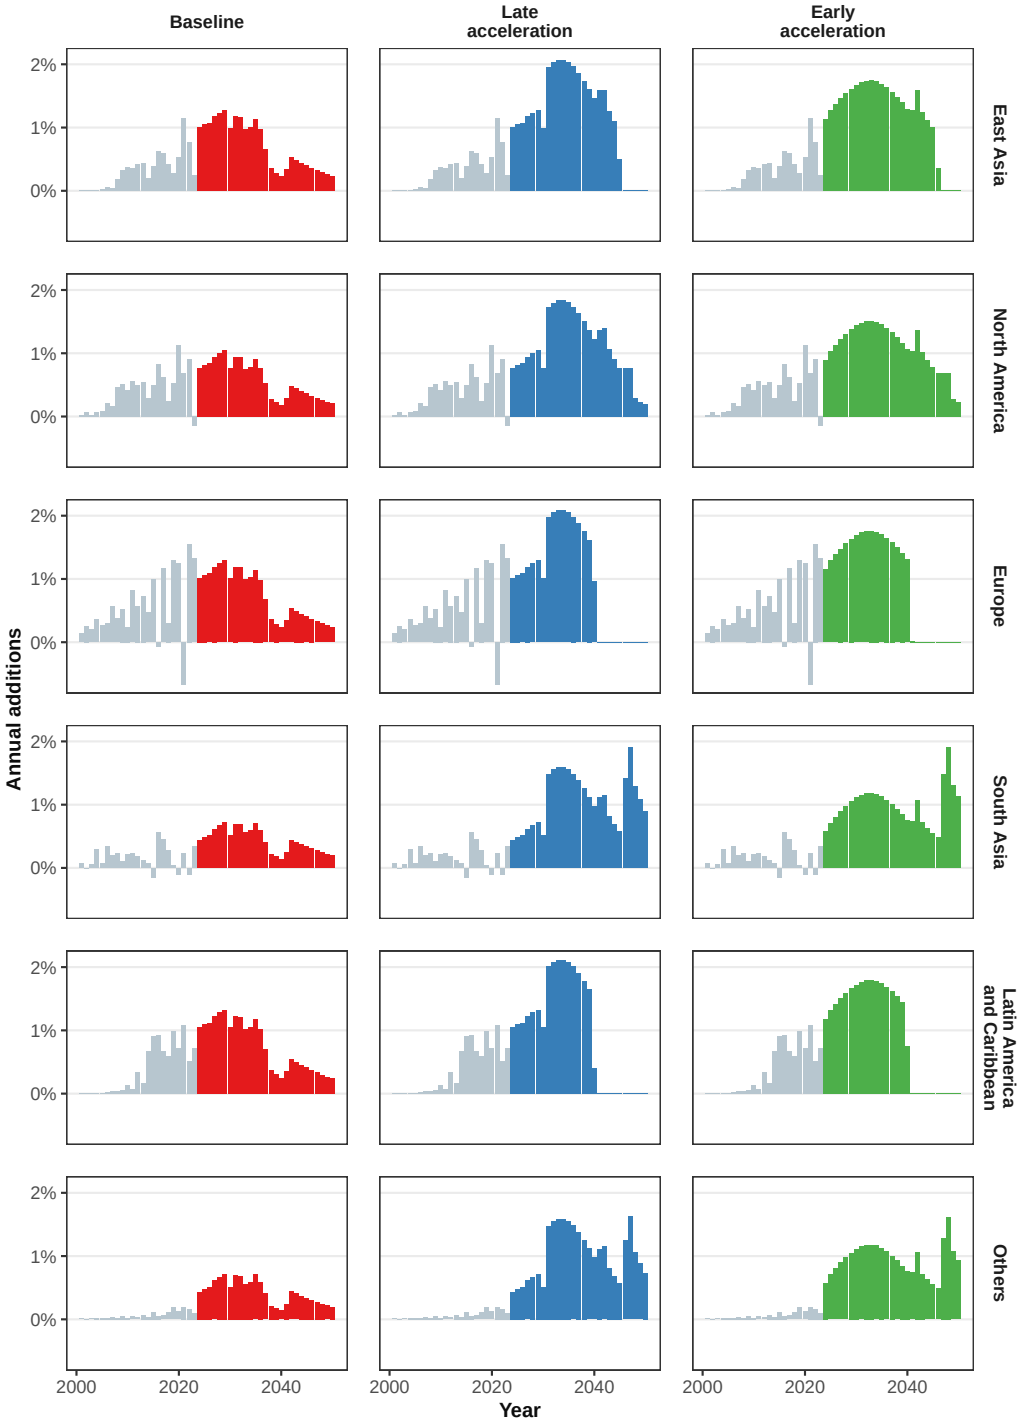

Supplementary Figure 13 Region-wise solar PV deployment under the Baseline, Early and Late acceleration scenarios.
Dots show historical data; solid lines show deployment in the Basline (red), Early acceleration (green) and Late acceleration scenarios (blue).

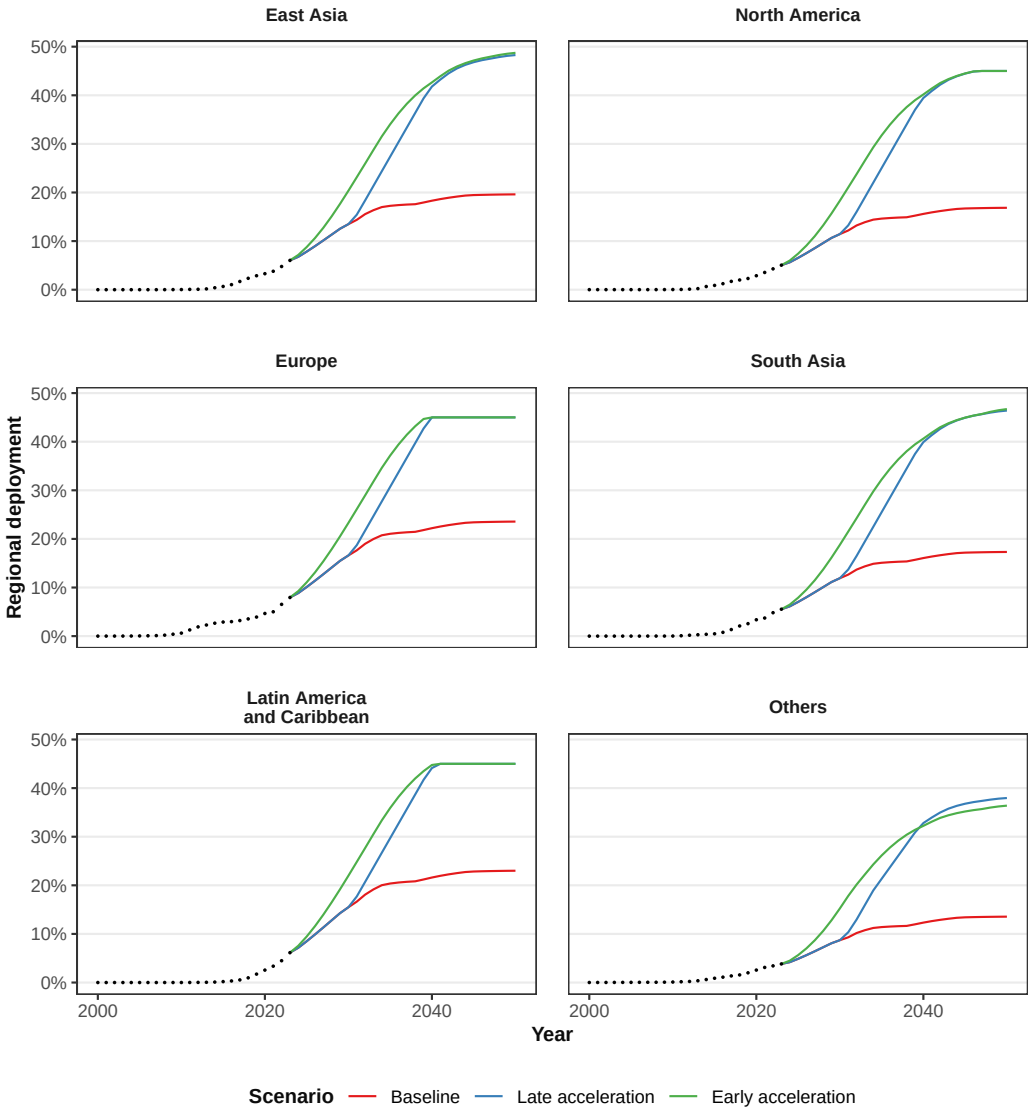

Supplementary Figure 14 Region-wise annual solar PV additions under the Baseline, Early and Late acceleration scenarios.
Gray bars show historical data; coloured bars show deployment in the Basline (red), Early acceleration (green) and Late acceleration scenarios (blue).

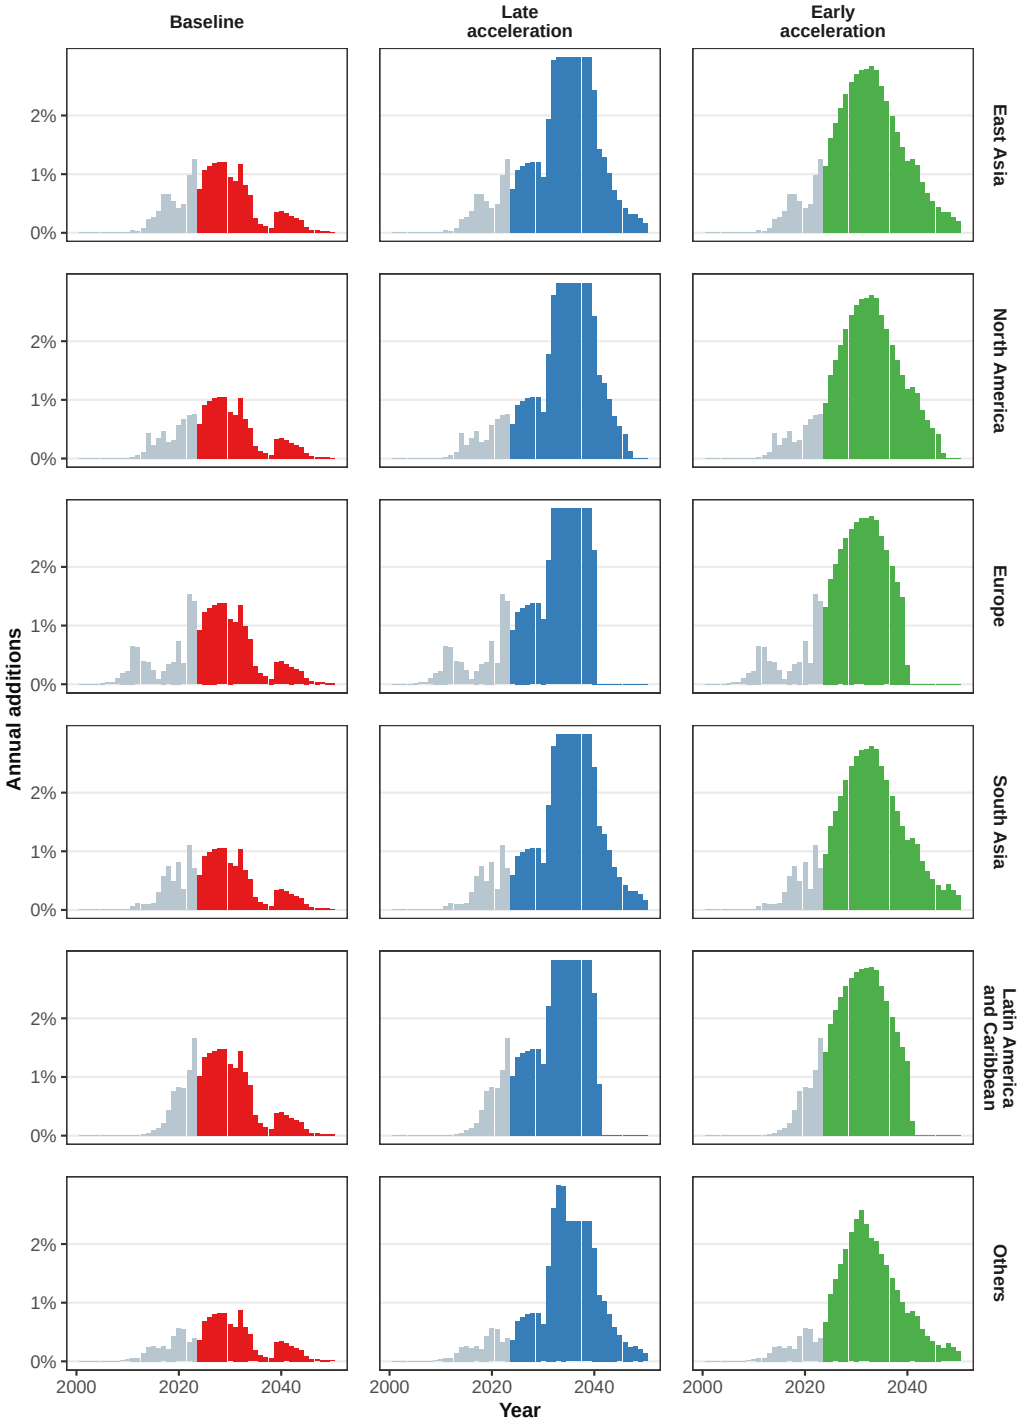

Supplementary Figure 15 Offshore wind deployment
Solid lines show historical deployment; stars indicate technology takeoff.

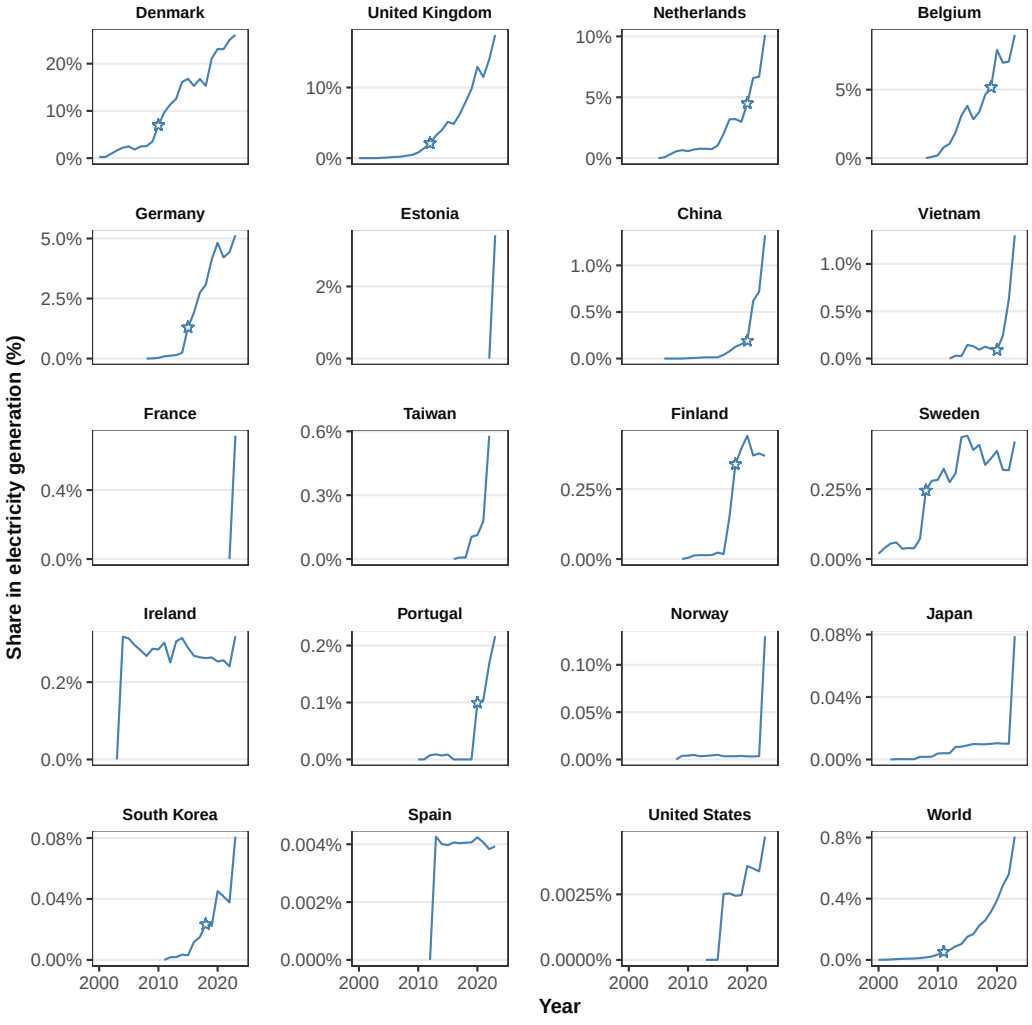

Supplementary Figure 16 Offshore wind rates
Bars show annual additions, solid lines show their three-year trailing average.

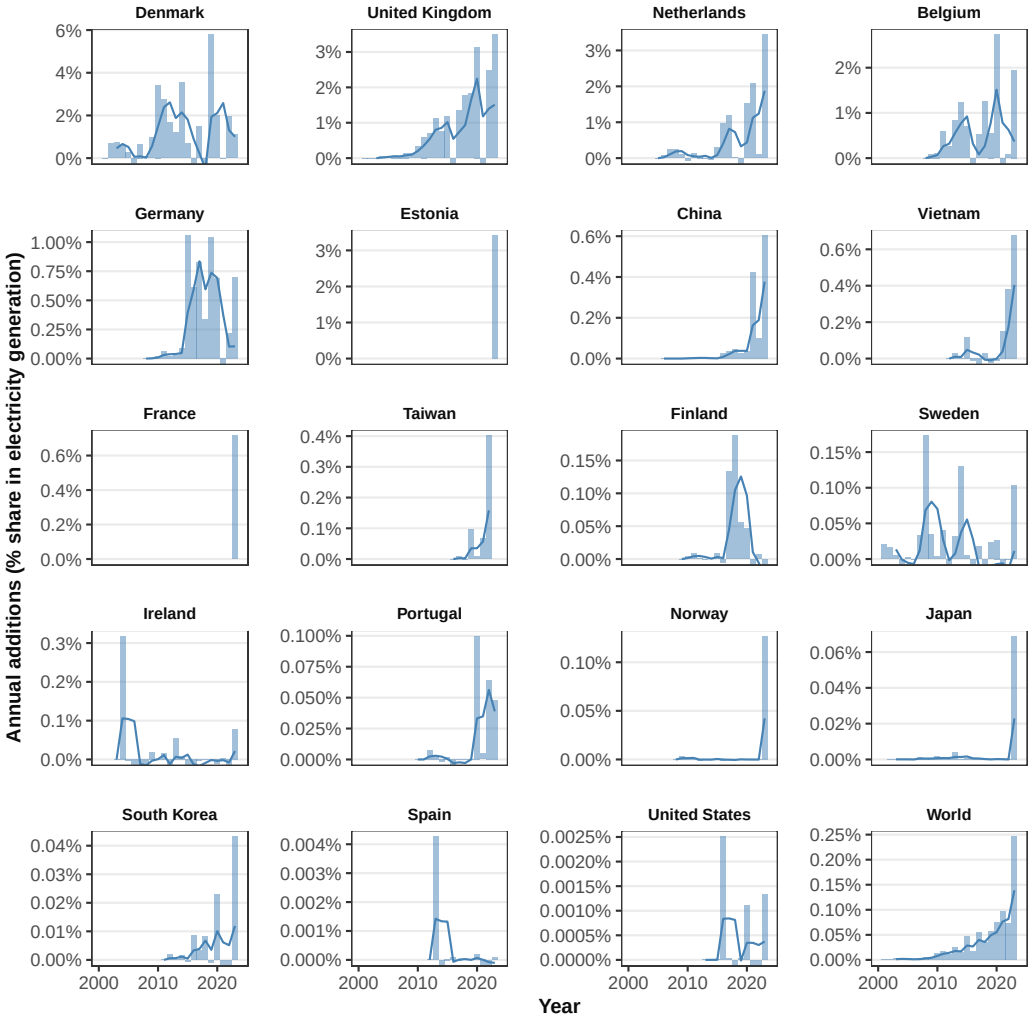

Supplementary Tables

Supplementary Table 1 Takeoff years by country for onshore wind.

Country	TO.Year	Cherp et al. (2021)	Country	TO.Year	Cherp et al. (2021)
DK	1991	1989	FO	2014	-
DE	1998	2000	PE	2014	2015
NL	1999	2003	JO	2015	-
ES	2000	1999	UY	2015	-
GR	2000	2001	ZA	2015	2015
JP	2001	-	LT	2016	-
NZ	2003	2005	PK	2016	2016
AT	2004	2004	FK	2017	-
IN	2004	2006	IL	2017	-
PT	2004	2003	NO	2017	2011
AU	2005	2007	IR	2018	-
CA	2005	2009	LV	2018	-
IE	2005	2000	AR	2019	2019
IT	2005	2007	GP	2019	-
KR	2005	-	KE	2019	-
FR	2006	2008	RU	2019	-
NC	2006	-	CU	2020	-
US	2006	2008	KW	2020	-
GB	2007	2007	SK	2020	-
BE	2008	2009	TW	2020	-
CN	2008	2010	VN	2020	-
SE	2008	2008	MD	2021	-
PL	2009	2010	MQ	2021	-
TN	2009	-	OM	2021	-
TR	2009	2010	SA	2022	-
MX	2010	2012			
CH	2011	-			
HR	2011	-			
BR	2012	2013			
CV	2012	-			
LK	2012	-			
LU	2013	-			
MA	2013	-			
UA	2013	-			
CL	2014	2014			
CR	2014	-			
EG	2014	2010			
FI	2014	2014			

Supplementary Table 2 Takeoff years by country for solar PV.

Country	TO.Year	Cherp et al. 2021	Country	TO.Year	Cherp et al. 2021
DE	2008	2009	RU	2020	-
ES	2008	2009	SA	2020	-
PT	2009	2014	SD	2020	-
AT	2010	2014	SV	2020	-
AU	2010	2013	VN	2020	-
IT	2010	2011	AM	2021	-
CA	2011	-	BA	2021	-
FR	2011	2014	BE	2021	2011
IN	2011	2016	BG	2021	2012
CH	2012	2014	BJ	2021	-
GB	2012	2014	BR	2021	2019
JP	2012	2013	BZ	2021	-
TH	2012	2015	CD	2021	-
DK	2013	2013	CL	2021	2015
KR	2013	2017	CO	2021	-
LU	2013	-	GR	2021	2012
US	2013	2016	HK	2021	-
CN	2014	2016	HR	2021	-
ER	2014	-	KE	2021	-
NL	2014	2016	LT	2021	-
ZA	2014	2016	MK	2021	-
FI	2017	-	MZ	2021	-
AW	2018	-	PK	2021	-
BF	2018	-	PL	2021	-
IL	2018	2014	RO	2021	2014
PE	2018	2018	RS	2021	-
PR	2018	-	SE	2021	-
AE	2019	-	SG	2021	-
AR	2019	-	SI	2021	-
CU	2019	-	SL	2021	-
CV	2019	-	TN	2021	-
HU	2019	2018	UZ	2021	-
LK	2019	-	AL	2022	-
MN	2019	-	ZM	2022	-
MX	2019	2019			
MY	2019	-			
SR	2019	-			
TW	2019	-			
UA	2019	-			
ZW	2019	-			
AF	2020	-			
AG	2020	-			
AO	2020	-			
BD	2020	-			
CI	2020	-			
CY	2020	-			
DO	2020	-			
EE	2020	-			
EG	2020	-			
ID	2020	-			
KH	2020	-			
LC	2020	-			
ML	2020	-			
NG	2020	-			
NZ	2020	-			
QA	2020	-			

Supplementary Table 3 Parameter configurations for onshore wind Monte Carlo simulations.

Config	Growth rate (k)	Ceiling (L)	Takeoff year	Design type
Base distribution parameters				
Base	$\Gamma(4, 8.163)$	$\Gamma(3, 39.474)$	$\mathcal{N}(5.5, 3.0)$	–
Grid exploration configurations				
1	0.144	0.157	10.5	Grid (low k, low L)
2	0.331	0.157	10.5	Grid (mid k, low L)
3	0.519	0.157	10.5	Grid (high k, low L)
4	0.144	0.319	10.5	Grid (low k, mid L)
5	0.331	0.319	10.5	Grid (mid k, mid L)
6	0.519	0.319	10.5	Grid (high k, mid L)
7	0.144	0.481	10.5	Grid (low k, high L)
8	0.331	0.481	10.5	Grid (mid k, high L)
9	0.519	0.481	10.5	Grid (high k, high L)
Face-centered configurations				
10	0.238	0.238	10.5	Face-centered (lower diagonal)
11	0.238	0.400	10.5	Face-centered (upper row)
12	0.425	0.238	10.5	Face-centered (right column)
13	0.425	0.400	10.5	Face-centered (upper diagonal)

Note: Base distributions are bounded ($0.01 \leq L \leq 0.75$, $0.1 \leq k \leq 1.2$). Parameter distributions are shifted to the means shown while maintaining the shape of the base distribution.

Supplementary Table 4 Parameter configurations for solar PV Monte Carlo simulations

Config	Growth rate (k)	Ceiling (L)	Takeoff year	Design type
Base distribution parameters				
Base	$\Gamma(4, 5.181)$	$\Gamma(3, 74.074)$	$\mathcal{N}(5.4, 1.98)$	–
Grid exploration configurations				
1	0.225	0.160	10.4	Grid (low k, low L)
2	0.521	0.160	10.4	Grid (mid k, low L)
3	0.817	0.160	10.4	Grid (high k, low L)
4	0.225	0.381	10.4	Grid (low k, mid L)
5	0.521	0.381	10.4	Grid (mid k, mid L)
6	0.817	0.381	10.4	Grid (high k, mid L)
7	0.225	0.602	10.4	Grid (low k, high L)
8	0.521	0.602	10.4	Grid (mid k, high L)
9	0.817	0.602	10.4	Grid (high k, high L)
Face-centered configurations				
10	0.373	0.271	10.4	Face-centered (lower diagonal)
11	0.373	0.492	10.4	Face-centered (upper row)
12	0.669	0.271	10.4	Face-centered (right column)
13	0.669	0.492	10.4	Face-centered (upper diagonal)

Note: Base distributions are bounded ($0.01 \leq L \leq 0.95$, $0.1 \leq k \leq 1.2$). Parameter distributions are shifted to the means shown while maintaining the shape of the base distribution.

Supplementary Table 5 Parameter configurations for mobile phones Monte Carlo simulations

Config	Growth rate (k)	Ceiling (L)	Takeoff year	Design type
Base distribution parameters				
Base	$\Gamma(4, 5.096)$	$\Gamma(3, 6.508)$	$\mathcal{N}(5.92, 2.68)$	–
Grid exploration configurations				
1	0.229	0.634	10.92	Grid (low k, low L)
2	0.530	0.634	10.92	Grid (mid k, low L)
3	0.831	0.634	10.92	Grid (high k, low L)
4	0.229	0.981	10.92	Grid (low k, mid L)
5	0.530	0.981	10.92	Grid (mid k, mid L)
6	0.831	0.981	10.92	Grid (high k, mid L)
7	0.229	1.327	10.92	Grid (low k, high L)
8	0.530	1.327	10.92	Grid (mid k, high L)
9	0.831	1.327	10.92	Grid (high k, high L)
Face-centered configurations				
10	0.379	0.807	10.92	Face-centered (lower diagonal)
11	0.379	1.154	10.92	Face-centered (upper row)
12	0.680	0.807	10.92	Face-centered (right column)
13	0.680	1.154	10.92	Face-centered (upper diagonal)

Note: Base distributions are bounded ($0.10 \leq L \leq 2.0$, $0.1 \leq k \leq 1.2$). Parameter distributions are shifted to the means shown while maintaining the shape of the base distribution.

Supplementary Table 6 Parameter configurations for CCGT Monte Carlo simulations

Config	Growth rate (k)	Ceiling (L)	Takeoff year	Design type
Base distribution parameters				
Base	$\Gamma(4, 9.091)$	$\Gamma(3, 7.752)$	$\mathcal{N}(6.0, 5.0)$	–
Grid exploration configurations				
1	0.133	0.448	11.0	Grid (low k, low L)
2	0.300	0.448	11.0	Grid (mid k, low L)
3	0.467	0.448	11.0	Grid (high k, low L)
4	0.133	0.569	11.0	Grid (low k, mid L)
5	0.300	0.569	11.0	Grid (mid k, mid L)
6	0.467	0.569	11.0	Grid (high k, mid L)
7	0.133	0.690	11.0	Grid (low k, high L)
8	0.300	0.690	11.0	Grid (mid k, high L)
9	0.467	0.690	11.0	Grid (high k, high L)
Face-centered configurations				
10	0.217	0.508	11.0	Face-centered (lower diagonal)
11	0.217	0.629	11.0	Face-centered (upper row)
12	0.383	0.508	11.0	Face-centered (right column)
13	0.383	0.629	11.0	Face-centered (upper diagonal)

Note: All configurations use a mixture of logistic (50%) and bilogistic (50%) models. Base distributions are bounded ($0.01 \leq L \leq 1.0$, $0.1 \leq k \leq 1.2$).

Supplementary Table 7 Regional deployment ceilings for onshore wind and solar PV in the acceleration scenarios.

Region	Onshore wind ceiling (%)	Solar PV ceiling (%)
East Asia	40	50
North America	40	45
South Asia	35	55
Europe	40	45
Asia-Pacific Developed	30	30
South-East Asia and developing Pacific	35	55
Africa	40	55
Eurasia	50	40
Latin America and Caribbean	35	45
Middle East	40	55

Note: Ceiling values represent maximum deployment potential as percentage of electricity generation.

Supplementary Table 8 Onshore wind deployment by acceleration scenario and region.

Scenario	Region	Final Deployment (%)	Peak Rate (p.p./yr)	Total Growth (p.p.)	Growth Share (%)
Baseline	*Global*	23.6	–	16.6	–
	Eastern Asia	26.8	1.28	–	41
	Europe	33.3	1.29	–	15
	Latin America and Caribbean	29.0	1.33	–	7
	North America	25.2	1.04	–	16
	South Asia	15.7	0.72	–	5
	Others	13.3	0.72	–	15
Early acceleration	*Global*	36.0	–	29.0	–
	Eastern Asia	40.0	1.74	–	40
	Europe	40.0	1.76	–	11
	Latin America and Caribbean	35.0	1.79	–	5
	North America	39.1	1.51	–	17
	South Asia	30.4	1.90	–	7
	Others	27.2	1.61	–	19
Late acceleration	*Global*	36.8	–	29.7	–
	Eastern Asia	40.0	2.07	–	39
	Europe	40.0	2.09	–	11
	Latin America and Caribbean	35.0	2.12	–	5
	North America	39.7	1.84	–	17
	South Asia	32.8	1.90	–	7
	Others	29.4	1.62	–	20

Note: Final Deployment and Total Growth shown as percentage of electricity generation; Peak Rate shown as annual percentage point increase (p.p./yr); Growth Share represents regional contribution to total global growth between 2023-2050.

Supplementary Table 9 Solar PV deployment by acceleration scenario and region.

Scenario	Region	Final Deployment (%)	Peak Rate (p.p./yr)	Total Growth (p.p.)	Growth Share (%)
Baseline	*Global*	18.3	–	12.7	–
	East Asia	19.6	1.20	–	38
	Europe	23.6	1.37	–	15
	Latin America and Caribbean	23.0	1.47	–	8
	North America	16.9	1.05	–	15
	South Asia	17.3	1.05	–	7
	Others	13.6	0.88	–	17
Early acceleration	*Global*	44.6	–	39.0	–
	East Asia	48.7	2.83	–	39
	Europe	45.0	2.87	–	12
	Latin America and Caribbean	45.0	2.88	–	6
	North America	45.0	2.79	–	17
	South Asia	46.7	2.79	–	8
	Others	36.4	2.57	–	18
Late acceleration	*Global*	44.7	–	39.1	–
	East Asia	48.3	3.00	–	38
	Europe	45.0	3.00	–	12
	Latin America and Caribbean	45.0	3.00	–	6
	North America	45.0	3.00	–	17
	South Asia	46.4	3.00	–	8
	Others	38.0	3.00	–	19

Note: Final Deployment and Total Growth shown as percentage of electricity generation; Peak Rate shown as annual percentage point increase (p.p./yr); Growth Share represents regional contribution to total global growth between 2023-2050.

Supplementary Notes

Supplementary Note 1: Technology growth and diffusion mechanisms and phases

The growth of new technologies is often represented as following an S-curve, where growth first accelerates, reaches a peak rate, and subsequently slows down and ends [6–8]. This mirrors the use of S-curves in disciplines like ecology and epidemiology and reflects the intuitive idea of early growth being driven by positive feedbacks before encountering increasing barriers and limits which ultimately prove unsurmountable.

The formative phase and technology take-off

Every new technology begins in the formative phase where its use is characterised by low deployment levels and erratic, unpredictable growth within a small set of niches [9–13]. Here, the growth and evolution of the technology is shaped by innovation, experimentation, and failures and as it "is tested, refined and adapted to market conditions" [14] with little interaction with the wider system. Together with the emergence of actor networks, business models and regulatory and policy environments, these developments allow the socio-technical regimes around the technology to facilitate consistent growth [10, 11, 15, 16].

This growth 'take-off' [11, 13, 17, 18] or "beginning of the movement" [6] signals the end of the formative phase and marks the first turning point in the technology's adoption. The existence of the formative phase has been documented at both the global [13] and country levels [11, 18]. At the global level it involves technology learning to reduce costs, standardisation of designs and establishing global supply chains, while at the country level it involves establishing necessary regulatory and market conditions suited to a specific national context.

The differences in when and which countries experience take-off illustrate the characteristics of the spatial diffusion of the technology [11]. However, while the literature proposes indicative thresholds for take-off – 2.5% [13] or between 0.3-1.8% [9] of the market, 100 MW installed capacity [19], 1% of total electricity supply [11, 20] – no take-off thresholds have been empirically established as yet. This presents a important challenge with regards to assessments of the role that emerging technologies might play in the future energy system. A technology that is beyond the formative phase in multiple markets may be judged to hold greater, more evidence-backed promise compared to one which isn't, but how do we make these assessments in the absence of a way to reliably measure when the formative phase ends? This also has implications for which policies are designed and implemented to accelerate

technological change. Supportive policies can play an important role in shortening the duration of the formative phase and enable earlier take-off within countries as well as faster cross-country diffusion.

Our contribution to this puzzle is two-fold: we first develop a method to empirically identify when different technologies take-off in different countries (Methods), and then use these estimates (Supplementary Figure 1) to measure how quickly technologies diffuse across countries (Supplementary Figure 2).

To measure this 'diffusion duration' we fit a logistic curve to the cumulative count for the number of countries where a technology has taken off and calculate $\Delta T_{diffusion} = ln(81)/k$, where k is the estimated growth constant. This metric indicates the years elapsed between the technology reaching take-off in 10% and 90% of countries. We show that while this cross-national diffusion took over half a century for older energy technologies like coal and hydropower, more recent (but still complex [21]) technologies like combined cycle gas turbines (CCGTs) and nuclear power have had shorter diffusion of durations of 36 and 30 years respectively. On the other end of the spectrum, mobile telephones, a granular [22], consumer-driven technology has had a diffusion duration of only 12 years. Solar PV and onshore wind fall in-between these cases with diffusion durations of 15 and 23 years respectively (Supplementary Figure 2).

The acceleration phase

After take-off, the technology enters the acceleration phase where it's growth begins to follow an S-curve. The acceleration phase is characterised by a consistent increase in annual deployment additions driven by positive feedbacks from technology and policy learning, and increasing investment profitability. These 'increasing returns' [23] or 'cumulative causation' [24] may manifest as an increase in the profitability of installing and operating new technologies or political gains from extending support to them at the national level. At the global level they may result from increasing economies of scale for manufacturing and servicing artifacts for the new technology and an expansion in the number of countries adopting the technology.

Consistent with S-curve behaviour, technologies in the acceleration phase experience a consistent decline in the year-on-year relative growth rates even as the annual additions are expanding in absolute terms. The annual additions keep increasing until they reach their maximum rate – this marks the second turning point in the technology's adoption where the positive feedbacks driving accelerating growth are balanced out by countervailing barriers [11].

These barriers can take various shapes and forms and appear at different stages of the adoption process, with their cumulative impact intensifying with increasing deployment. They can include

resistance from incumbents, declining social acceptance and increasing public opposition, system integration challenges, land and other resource-use constraints, limits to supporting infrastructure, supply chain congestion, or limits to institutional capacity [25–38].

Growth pulses and the stable growth phase

It is typically assumed that a technology's growth immediately slows and approaches the saturation phase after it achieves its peak growth rate as the increasing barriers to adoption overpower the mechanisms driving it. At the end of the saturation the technology's growth grinds to a complete halt and stagnates at its peak market share or final deployment ceiling [39]. Following this point, the technology may maintain its market share until it begins to lose ground to an emerging alternative.

Contrary to this idea, we find that in some cases, the technology may experience a fresh 'pulse' of growth due to the availability of new markets, or the (re-)introduction of supportive measures. These growth pulses can be observed for several technologies at both the national and global scales. At the national scale, these pulses may be linked to the introduction of new policies, changing regulations, market reforms, financial support schemes, or due to other changes in the socio-technical landscape. At the global scale, they may result from market expansion as a technology diffuses to a new set of countries, or renewed growth in countries where growth had previously begun to slow or stagnate. For some technologies including nuclear power, coal and hydropower we see a renewed pulse of global additions due to delayed but large-scale adoption in Asia with China driving a substantial share of new growth (Supplementary Figure 4).

Together, the interplay of these mechanisms often culminates in a prolonged stable growth phase, with a balance between drivers and barriers yielding a period of nearly-linear growth. The presence of this stable phase hints at the role that targeted policy effort can play in sustaining growth over longer periods of time and delaying saturation. Policy effort can induce diffusion through cross-country technology transfers and knowledge spillovers, accelerate takeoff by helping curtail the duration of the formative phase, prolong the stable growth phase by introducing measures that address emerging barriers and delay saturation, and catalyse renewed growth pulses by creating new incentive structures.

In sum, we argue that the growth of a new technology proceeds through a sequence of four distinct phases – formative, acceleration, stable growth, and saturation. These phases are repeated in every country at different points in time, and global deployment patterns emerge from the aggregate effects of national adoption. Initially global growth is largely influenced by technology take-off and acceleration in a small number of early adopters. Subsequently, it is shaped by asynchronicities

in cross-national diffusion, takeoff, acceleration, peaking and saturation which generate a myriad co-evolving feedbacks, synergies and dissonance.

The measurement of technology growth

Our analysis also shows that there are different metrics which can be used to measure technology adoption, and that the choice of metric influences the inferences one draws about the nature of a technology's growth. We measure technology growth in relation to the larger market – shares in electricity generation for solar and wind for which we have better generation data, shares in installed electricity generation capacity for other energy technologies, and share in population for mobiles.

Technology adoption is also often measured in absolute units (e.g. GW of solar capacity or MWh of solar electricity generation), but this ignores the crucial relationship between a technology and the market it operates in – when we measure adoption, a technology's market share is not only influenced by its own deployment growth dynamics but also those of the larger market within which it operates.

We show that the earliest electricity generating technologies based on coal, oil and hydropower were often responsible for creating a new market for electricity when they first began to be deployed. This led to their market shares starting at close to 100% and then declining when other technologies entered the market (Supplementary Figure 3). Their deployment patterns are markedly different from later technologies like nuclear power that entered a pre-existing market for electricity and saw their shares gradually increase from zero. The example of nuclear power also illustrates how the size of the overall market impacts the characteristics of the technology's adoption curve – because of the relatively smaller size of the overall electricity system at the time of its adoption, even small increases in absolute deployment could translate to relatively big jumps in the technology's market share.

Thus, for many historical technologies, focusing on absolute versus relative adoption metrics would lead to significantly different inferences about the nature of growth. For more recent technologies like CCGTs, mobiles, solar PV and onshore wind, we see a near-perfect alignment between the patterns for the two, as absolute deployment went hand-in-hand with increasing market penetration amid relatively modest changes in the size of the overall market.

Supplementary Note 2: Mathematical models for technology growth

Technology S-curves are frequently represented using the logistic model [40], a three parameter function of the form

$$f(t) = \frac{L}{(1 + e^{-k(t-t_0)})} \tag{1}$$

, with parameters k (the growth constant), L (the final deployment asymptote or 'ceiling'), and t_0 (the inflection point). Deployment achieves its maximum annual growth rate at the inflection point, which can be expressed as $G = \frac{Lk}{4}$ [11].

Logistic curves fit to empirical deployment data have been widely used for ex-post description and comparison of growth for historical technologies which have already reached their final ceilings[6–8, 41]. Such analyses are less suited to emerging technologies with expanding deployment due to the limited reliability of parameters estimated by fitting logistic functions to time-series with continuing growth [11, 42, 43]

We posit that different parameters of the logistic curve reflect information about mechanisms shaping growth at different phases of the S-curve: the growth constant k captures the rate of deployment expansion during the acceleration phase, the inflection point t_0 marks the start of the quasi-linear stable growth phase, and the deployment ceiling L indicates where growth finishes. Parameters reflecting later phases cannot be reliably estimated from early observations because these data are unable to represent all phases of technology growth – as they are only able to capture dynamics from a limited set of early phases, they do not have any information about the evolving mechanisms that shape later phases of growth. Thus, fitting a logistic function to data from the acceleration phase with consistently increasing annual additions might give us a reasonable estimate for k, but will only be able to guess at the values of t_0, or L as the underlying data does not have any information on the phases to come. The more data we have and the better its coverage of the whole deployment curve, the more reliable our parameter estimates.

We illustrate this phenomenon by fitting a logistic function to systematically truncated deployment timeseries for different technologies. For each technology, we track the evolution of the parameters estimated by fitting the curve to curtailed data ending in a specific year. Our analysis shows that while the estimate for k stabilises fairly early (essentially as soon as the technology is nearing the end of the acceleration phase), the estimate for L is extremely unstable until very close to true saturation (Supplementary Figure 6). This 'moving L' phenomenon (Supplementary Figure 5), where the estimate for the ceiling regularly changes with increasing deployment makes it particularly difficult to use logistic fits to project future technology growth and often leads to overfitting (when the curve describes the data it is fit to reasonably well but is unable to accurately predict future growth).

Several different approaches have attempted to address this parameter estimation challenge. We see these as falling into two broad categories – one focusing on modifying the standard logistic curve

fitting approach, and another looking beyond the logistic function and introducing alternative growth models.

In the first category, ref. [44] suggest using an exogenously assumed fixed value for L (aligning with a normative target or goal such as the 1.5ºC target) with a set take-off year and empirically derived estimates for the 'emergence rate' (which is analogous to the growth constant k) to generate a logistic curve. While the trajectories generated using this approach can be used to assess if growth is on track to reach the desired level by a certain time given a certain emergence rate, the projected is purely conditional on the choice of L and the emergence rate.

Another modelling framework from ref. [43] improves on this 'fixed L' approach by having L linearly increase with time at an exogenously define rate. This formulation more accurately captures the 'moving L' phenomenon we highlight, but is again sensitive to assumptions about how fast L changes and at what level it peaks. Thus, while it is useful for constructing feasibility spaces for future growth under different assumptions about the emergence rate and the 'demand pull' raising L, it does not resolve the fundamental problem of empirically deducing L from early data.

The second category focuses on exploring alternatives to the logistic model, and has had a significantly longer history than the first.

Ref. [45] introduced the Gompertz model [46] to analyses of technology diffusion already in 1980. The Gompertz model, an asymmetric S-curve of the form

$$f(t) = Le^{-e^{-k(t-t_0)}} \tag{2}$$

with parameters k (the growth constant), L (the final deployment ceiling), and t_0 (the inflection point), has a longer growth phase with delayed saturation which mitigates some of the logistic model's pessimism about the ceiling when used with earlier observations. However, it is also vulnerable to the same problem of being unable to reliably estimate parameters describing later phases from early data [11, 42].

Given the accelerating nature of growth early in the S-curve it is sometimes (implicitly [47–49] and explicitly [2, 43]) argued that the acceleration phase can be described by an exponential function. The exponential function is of the form

$$f(t) = y_i e^{K(t-t_i)} \tag{3}$$

where e is the exponential growth constant, and y_i is the value at initial time t_i. While it's unsuitability for describing or projecting technology growth over the long-term due to its inherent inability to account for any barriers and slow-down is quite obvious, we show that it is also ill-suited to describing the acceleration phase because it does not capture the declining year-on-year growth rates characterising it – by definition, the exponential model assumes a fixed growth rate equal to the growth constant.

Recent studies [50, 51] have demonstrated yet another approach which attempts to mitigate the deficiencies of individual models by using several different models and averaging projections from across the ensemble. They fit different growth models to empirical data, generate projections, use hindcasting to evaluate the performance of each model, and then generate weighted projections where better performing models are assigned higher weights. While the use of a more diverse ensemble of models coupled with iterative hindcasting makes this approach less vulnerable to overfitting compared to projections based on a single model fit to a single timeseries, it still faces the same underlying problem of estimating parameters corresponding to later phases using early data.

We argue that this parameter estimation problem is not limited to a particular growth model but extends to the broader enterprise of using models fit to empirical data to make projections.

Supplementary Note 3: PROLONG (PRobabilistic mOdeL Of techNology Growth)

Models fit to a single set of historical observations are implicitly constrained in their ability to anticipate future technology growth dynamics. A logistic function fit to solar PV data until 2015 does not reliably project global deployment in 2023 because it only captures dynamics from the acceleration phase. We face the same problem when we try and project deployment until 2030 or 2040 using the latest data for wind or solar PV because these technologies are still in the acceleration phase at the global scale. But this does not preclude the possibility of us making meaningful projections altogether as there are individual countries where these technologies are already at more advanced phases of adoption [11]. We argue that even though deployment data from these countries do not yet capture slow-down and saturation, they offer us an empirical window into studying the evolving balance of drivers and barriers across the acceleration and stable growth phases. Here, we develop an approach to use these 'incomplete' national deployment time-series to probabilistically project global growth over the near to medium-term.

We find that statistical ranges for the growth constant (k) and peak growth rate (G) from logistic curves fit to national data are relatively stable over time and can help anchor our expectations for the eventual values of these parameters for the world as a whole (Supplementary Figure 6). At the same time, our analysis also shows that similar statistics for L at the national scale are not informative for estimating L at the global scale. We posit that early national k and G from countries at more advanced phases of adoption could be used to inform projections of global growth. However, doing so requires addressing some important challenges.

First, we would need to be able to identify the mathematical relationship between national parameter statistics measured at earlier points in time and the parameters of a curve describing the completed global deployment trajectory. This presents the same problem contemporary approaches struggle with – in the absence of a crystal ball that shows us data from the future, we have no way of knowing what the parameters of the eventual global trajectory will be.

One way to overcome this challenge would be to study historical technologies for which we have data spanning the whole S-curve and use them as reference cases. We could quantify the relationships between early national curve parameters and final global curve parameters, and then use them with other technologies. But how do we know if the relationships measured for one reference technology also hold for another? Given the specificity of technological characteristics and the socio-political, economic context in which each technology is deployed, it is difficult to guarantee that patterns for one technology can be used to predict those for another. If we want to project the future deployment of solar power, is it likelier to follow patterns observed for other energy technologies like nuclear power or CCGTs during the 20th century, or for granular consumer technologies like mobiles in more recent years?

Another challenge relates to the actual process of quantifying these relationships between national and global parameters – what methods can we use to capture these complex, non-linear relationships that might vary in shape and form across different technologies?

We resolve these challenges by using computational simulations to explore different possibilities for the growth and diffusion of a given technology in an ensemble of virtual worlds, and using a machine learning model to capture the quantitative relationships between parameters describing incomplete national growth and completed global trajectories.

We argue that a diverse enough possibility space composed of thousands of simulated global trajectories, each of which is the aggregate outcome of a unique set of national growth dynamics, captures adoption patterns similar to those unfolding in the real world, as well as those where the technology

is more/less successful. In generating these simulations, we make explicit assumptions about what we see as a plausible range for the takeoff timing, for the shape and form of national growth trajectories, and for permissible deployment speeds and ceilings. Our approach uses computational simulations to train a machine learning model to recognise the relationships between early national growth parameter statistics and the parameters of a technology's final global trajectory, and then use the trained model to make probabilistic projections for global growth using empirical national data.

The implementation of our modelling framework consists of the following steps:

- **Step 0:** Defining rules for the virtual worlds
- **Step 1:** Exploring diverse technology futures using Monte Carlo simulations
- **Step 2:** Generating training data from the ensemble of simulated trajectories
- **Step 3**: Using machine learning to identify the relationships between curtailed national and final global parameters
- **Step 4:** Generating probabilistic projections from empirical data
- **Step 5:** Model validation and hindcasting

Step 0: Defining rules for the virtual worlds

Before we start generating simulations we define a set of core assumptions that form the basic structure for the model.

Our simulation module represents a virtual world composed of 150 countries, with the distribution of their relative, technology-specific market sizes carefully calibrated to reflect real-world market sizes. We quantify the market shares as normalised shares (summing to 1). For solar PV and onshore wind, market shares are derived from national total electricity generation, reflecting each country's relative electricity system size. For mobile phones, market shares are proportional to population counts, representing the potential user base. For CCGTs, market shares are calculated based on each country's existing natural gas electricity generation capacity. These market shares directly weight each country's contribution to global aggregate deployment. For example, a country with 5% of global electricity generation would contribute five times more to global solar deployment than a country with 1% generation share, assuming identical percentage-based deployment within each country. This market-weighted approach ensures our simulations properly account for the outsized influence of large markets on global technology diffusion patterns, while still capturing the diversity of growth dynamics across different national contexts.

We simulate technology growth in each country separately, and global deployment at each time step is the weighted sum of all national deployment at that instant.

We assume that each country can follow trajectories modelled by one of two shapes – a standard S-curve (using a logistic function) or a curve with two pulses (using a bi-logistic function). For the logistic trajectories, we need to assign each country a takeoff year (the year when deployment first exceeds 1%), intrinsic growth rate (k) and deployment ceiling (L), after which we can calculate the inflection point using the equation $t_0 = takeoff_{year} + (1/k)log(L/0.01 - 1)$, where t_0 is the inflection point year. For the bi-logistic trajectories, we need to assign each country a takeoff year, intrinsic growth rate for the first pulse (k_1) and the final deployment ceiling (L). The final ceiling is split between the two pulses with respective ceilings L_1 and L_2 using a random fraction between 20-80%. We calculate the inflection point for the first pulse (t_{0_1}) and its peak annual growth rate (G_1). The second pulse starts after a random delay between 1-10 years after t_{0_1}, and is assigned a peak annual growth rate (G_2) calculated by multiplying G_1 by a random multiplier that is centred around 1 – this allows for the equal possibility of the second pulse exhibiting faster/slower growth than the first. The intrinsic growth rate (k_2) for the second pulse is calculated as $k = 4G2/L2$, and its inflection point (t_{0_2}) as $t_{0_2} = takeoff_{year_2} + (1/k2)log(L2/0.01 - 1)$. The choice of a 1-10 year delay for the second pulse allows for both quick transitions between pulses (1 year) and longer pauses between growth phases (up to 10 years).

Real-world technology deployment rarely follows perfect (bi-)logistic curves due to policy changes, economic fluctuations, supply chain disruptions, and other temporal factors. To replicate this variability, we implement a correlated noise model with three parameters. First, we apply multiplicative noise with an initial amplitude of 5% of the current deployment value. This ensures deviations scale proportionally with deployment levels – smaller variations in early adoption phases and larger absolute fluctuations during rapid growth periods. Second, we implement year-to-year correlation ($\rho = 0.7$) to create persistent effects that mimic how real-world drivers and barriers typically influence deployment over multiple consecutive years. For instance, the implementation of a new policy might affect deployment over several years rather than causing independent annual fluctuations. Third, we incorporate amplitude decay (2% annual reduction from the initial 5%) to reflect how mature markets tend to demonstrate more stable, predictable growth with reduced volatility. This three-component noise structure creates trajectories with realistic short-term variability while preserving the underlying diffusion pattern, avoiding both the artificial smoothness of perfect curves and the unrealistic randomness of uncorrelated noise.

Here is an example to illustrate how the noise works. Let's consider a country with 10% deployment in year t. With a 5% noise amplitude, the deployment value might be adjusted to $10\%(1+0.05) = 10.5\%$ or $10\%(1-0.05) = 9.5\%$, depending on the random noise value. In year $t+1$, if deployment grows to 15% and the previous year's noise was +5%, the correlated noise would be calculated as $0.7(+5\%) + 0.3$ (the new random noise). If the new random component is -3%, then the noise in year $t+1$ would be $0.7(+5\%) + 0.3(-3\%) = +2.6\%$. The deployment would then be adjusted to$15\%(1 + 0.026) = 15.39\%$. By year $t + 5$, the noise amplitude would have decayed to $5\%(0.98)^5 \approx 4.5\%$, reducing the magnitude of potential fluctuations as the market matures.

We introduce spatial heterogeneity in when growth starts, how fast it happens, and where it ends by assigning countries different parameters from different statistical distributions such as normal or gamma distributions which vary between technologies. We use a Monte Carlo simulation engine to simulate technology growth in thousands of virtual worlds, where each country follows a new, unique trajectory every single time, leading to the emergence of a diverse set of global growth patterns. The Monte Carlo simulation is implemented using R's parallel processing capabilities through the 'parallel', 'foreach', and 'doParallel' packages, allowing for efficient computation across multiple CPU cores simultaneously.

We simulate technology growth and diffusion over a time period of 50 years; long enough for countries to move from the formative to slow-down phases.

Step 1: Exploring diverse technology futures using Monte Carlo simulations

Establishing parameter distributions

To start generating simulations, we first define the distributions which control the national growth dynamics.

We first look at early empirical data for the technology – up to 2015 for solar PV, 2010 for onshore wind, 1998 for mobiles, 1985 for CCGTs – and fit normal distributions to the available national takeoff years, and gamma distributions to logistic k and logistic L for mature countries (with logistic curve maturity $\geq 50\%$). These distributions inform our priors about national curve parameters for the specific technology at an early stage of adoption and define our 'base configuration'. Given what we know from empirical analyses of technology growth, these distributions most certainly do not capture the 'true' parameter space for the eventual growth of the technology and very likely underestimate its full potential. The gamma distribution is chosen for modeling k and L parameters because it naturally constrains values to be positive while allowing for right-skewed patterns.

We address this limitation by first defining a broader range of plausible values for the takeoff year, k and L for each technology (with the specific aim of accounting for a diverse set of diffusion patterns) and then systematically shifting the distributions for each parameter to explore different regions of the technology's possibility space. We use a structured approach that ensures thorough coverage of parameter combinations—creating worlds where technologies grow quickly but saturate at low levels, others where growth is slower but reaches higher penetration levels, and various combinations in between. For example, one configuration might combine higher k values with moderate L, while another might pair slower k with higher L (Supplementary Tables 3-6).

We implement this parameterization using a hybrid coverage approach that combines a full factorial design with face-centered points to efficiently explore the parameter space. Our code creates a grid of parameter values by dividing each parameter range (k, L, and takeoff year) into multiple divisions and selecting points at these positions. For example, with $k_{divisions} = 3$, we sample 3 different mean values for the gamma distribution of k, from lower to higher growth rates. Consider two adjacent grid points in our parameter space: one with $k_{mean} = 0.2$, $L_{mean} = 0.4$ and another with $k_{mean} = 0.3$, $L_{mean} = 0.6$. The face-centered approach adds an additional point between them at $k_{mean} = 0.25$, $L_{mean} = 0.5$. This ensures we capture not just corner cases but also intermediate parameter combinations, providing more comprehensive coverage of the parameter space. For each configuration, we adjust the underlying distribution parameters (shape and rate for gamma distributions, mean and standard deviation for normal distributions) to achieve the target mean values while maintaining appropriate dispersion. This approach generates 13 distinct configurations that systematically explore different combinations of early/late takeoff, fast/slow growth, and low/high saturation levels.

Running Monte Carlo simulations

We generate 13 different configurations and then use each one to simulate 1000 different technology diffusion pathways. For each configuration, we run a Monte Carlo simulation engine which randomly assigns parameter values to individual countries by drawing from the distributions defined in this configuration. Each country is assigned its own takeoff year, intrinsic growth rate (k) and deployment ceiling (L), which are then used to generate its deployment trajectory. We account for different market sizes by weighting each country's contribution proportionally to its share of the global market, simulating how technologies might spread differently in large versus small countries.

These country-level trajectories are then aggregated into a global adoption curve for each simulation run. By repeating this process 1000 times for each configuration, we generate a dataset of

13,000 plausible global diffusion trajectories—some showing rapid global adoption, others displaying more gradual growth patterns, and many exhibiting complex multi-phase growth.

In addition to exploring the possibility space by varying the parameter configurations, we also run these simulations under differing assumptions about the shape of national growth trajectories. We run three different batches of simulations – one where all countries have logistic growth, one where all countries have bi-logistic growth, and one where each country has an equal chance of following logistic or bi-logistic growth and there's an even split in the number of countries following either pattern in each run. Subsequent steps involving the generation of training data and producing a trained machine learning model are followed separately for each batch – see Step 4 for more details.

We can make our coverage of the possibility space more comprehensive by running a larger number of simulations and using more granularly differentiated configurations, but each addition comes at the cost of significantly higher computational requirements.

To make subsequent analysis more computationally tractable, we apply a strategic filtering process to select a smaller, more diverse subset for our final training dataset. Rather than randomly sampling or using all runs (which would be computationally expensive and potentially include redundant runs), we employ a clustering-based approach that identifies distinct pattern families within each configuration's simulation runs. The algorithm analyses the characteristics of each trajectory – such as k, inflection points, and final saturation levels – and groups similar trajectories together. From each cluster, we select representative runs that best capture that pattern family's essential features. This approach ensures our training dataset maintains the full diversity of possible diffusion patterns while substantially reducing its size – from 1000 total runs down to 200 representative runs per configuration. The selection process works as follows:

- For each run, we extract feature vectors that characterize the growth patterns, including growth parameters of large countries, correlations between market size and growth parameters, heterogeneity in growth rates, and leader-follower dynamics
- We normalize these features and apply principal component analysis (PCA) to reduce dimensionality while preserving approximately 85% of the variance
- We then apply k-means clustering (implemented in R's base 'kmeans' function) to group similar trajectories
- From each cluster, we select the medoid run (the run closest to the cluster center) as the representative for that pattern family.

The selection of 200 runs per configuration was determined through experiments balancing computational efficiency with maintaining diversity – we found this range provided sufficient pattern coverage while substantially reducing computation time for subsequent model training.

Step 2: Generating training data from the ensemble of simulated trajectories

For the final step in our data preparation, we transform our simulation data into a format that teaches the model how to make predictions with limited information. Here, we systematically truncate the national deployment data from each simulation run at various early years (years 12, 15, 18, 21, 24, and 30) to create snapshots of what the diffusion pattern would look like if observed at those points in time. This exercise mimics the real-world challenge of forecasting from partial historical data. For each curtailed snapshot, we fit logistic growth curves to the curtailed national-level data, extracting parameters that characterise the diffusion process in each country up to that point. These include median, Q1 and Q3 for k and G, as well as the number of "mature" countries in the stable growth phase and their combined market share. We only use observations from those countries where the logistic curve maturity is identified as being over 50

The parameter extraction process employs a curve-fitting procedure implemented adapted from ref. [11]. For each country with sufficient data points (at least 5) and meaningful deployment ($\geq 1\%$), we fit a logistic curve and extract the key parameters. We then calculate distributional statistics (median, first quartile, third quartile) across all mature countries to characterise the overall pattern of national growth. This approach captures both the central tendency and variation in growth patterns among early adopters.

We then pair these early-stage national parameters for each simulation run with the known parameters of a logistic curve fit to its full global trajectory, creating training data that connect "what we know so far" at the national level with "what eventually happens" at the global scale. We feed this data into a quantile random forest model – an ensemble learning method that builds hundreds of decision trees, each of which learns decision rules that connect early-stage national diffusion patterns to their ultimate global outcomes.

Step 3: Using machine learning to identify the relationships between curtailed national and final global parameters

We implement the quantile random forest using the 'ranger' package in R, a high-performance implementation of random forests particularly well-suited for large datasets. Our model uses 1000 trees,

which we found sufficient to achieve stable predictions while balancing computational efficiency. We enable the "quantreg" option in ranger to estimate conditional quantiles rather than just the conditional mean. The importance of features is calculated using the impurity-based method, allowing us to identify which national parameters most strongly influence global outcomes.

The forest's structure captures the complex, non-linear relationships between early diffusion signals and long-term outcomes, while being robust against outliers and noise. Unlike standard regression models that predict only mean outcomes, our quantile forest approach estimates the entire conditional distribution of possible futures.

We train two separate models – one that uses early national data to predict the global k, and another which uses the same data to predict the global G. By training on a diverse set of simulated histories with known outcomes, the models develop an understanding of how partial patterns tend to evolve, enabling them to look at real-world data up to the present day and make informed projections about the full spectrum of likely future evolutions—from conservative lower bounds to ambitious upper estimates.

Step 4: Generating probabilistic projections from empirical data

To use the model to make projections, we prepare input data capture the median, Q1 and Q3 k and G for mature logistic fits to empirical national data, the number of mature countries in the sample, their combined global market share, and the current global deployment level.

The model takes these inputs and predicts the quantiles for k and G, which are then used to derive a value for the saturation level (L). We consider all possible combinations of k and G quantiles (5x5=25 combinations), and filter out those combinations where $L \leq$ the current deployment. Each combination is then used to generate a complete logistic curve, with the inflection point calculated using the current global deployment (y_i) and current year (t_i) using the equation$t_0 = ((1/k) * log((L/y_i) - 1)) + t_i$, where t_0 is the inflection point year, k is the growth rate, L is the saturation level, y_i is the current deployment, and t_i is the current year.

For each valid k-G combination, we generate a full trajectory up to the desired projection horizon. The global growth constant (k) directly influences how quickly deployment accelerates and then decelerates, while the peak annual growth rate (G) helps determine the saturation level (L) through the relationship $L = 4G/k$. This approach captures the mathematical relationship between these parameters while ensuring consistency with current observed deployment levels.

To identify the most representative trajectory for each quantile level (5th, 25th, median, 75th, and 95th), we calculate the area difference between each generated trajectory and the year-by-year quantile values derived from all valid trajectories. The trajectory with the minimum area difference becomes the representative curve for that quantile, providing a coherent set of parameters that best represents that particular future pathway. This area-based matching ensures that the representative trajectories maintain consistent parameter relationships while closely tracking the statistical properties of the full ensemble.

For each year in the projection horizon, we collect the predicted deployment values from all valid trajectories and calculate summary statistics across the full set which gives us a year-by-year probability distribution for future global deployment.

Our probabilistic framework explicitly acknowledges multiple dimensions of uncertainty. First, the use of quantile random forests captures the uncertainty inherent in technology diffusion processes—the natural variability observed even when initial conditions are similar. Second, by generating projections from diverse parameter combinations rather than single parameter values, we address parametric uncertainty about the true values of growth parameters. Finally, by training separate models on different trajectory shapes (logistic, bilogistic, and mixed), we incorporate structural uncertainty about the underlying model form. This comprehensive treatment of uncertainty provides a more realistic view of possible futures than deterministic approaches or those that address only a single dimension of uncertainty.

Step 5: Model validation and hindcasting

To validate the predictive capabilities of our approach, we implement a comprehensive hindcasting framework that systematically evaluates how well our models can project known historical deployment patterns from earlier, incomplete data.

We perform a series of hindcasting tests to evaluate the performance of our projection model. We do so by using the model with empirical national data truncated at different years in the past, and then comparing the resulting projections to 'out-of-sample' global data.

Our hindcasting approach follows these steps:

1. For each technology, we select multiple historical cutoff years (e.g., 2005, 2010, and 2015 for solar PV)
2. For each cutoff year, we extract national deployment data up to that year only
3. We apply our model to this truncated data to generate projections

4. We compare these projections against the actual observed deployment in subsequent years
5. We repeat this process for models trained on logistic, bilogistic, and mixed trajectory data

This process mimics real-world forecasting scenarios and provides an objective assessment of each model variant's predictive performance.

To estimate the absolute error, we use the symmetric mean absolute percentage error (sMAPE) [52] which quantifies the absolute magnitude of point errors (thereby avoiding the cancellation of negative and positive errors) and also accounts for the relative scale of the quantity being measured (thereby avoiding asymmetricity). We use it to assess the overall out-of-sample forecasting performance for each model when informed by in-sample observations until year z:

$$sMAPE = \frac{1}{t_{end} - z} \sum_{i=z}^{t_{end}} \frac{|y_z - f(z)|}{(|y_z| + |f(z)|)/2} * 100 \tag{4}$$

To estimate the directional error, we calculate the symmetric mean percentage error (sMPE) [52] which quantifies both the magnitude and direction of the point errors and avoids asymmetricity and large errors (when the out-of-sample values are close to zero). We use it to assess the tendency of a model informed by in-sample observations until year z, to over- or under-predict out-of-sample deployment:

$$sMPE = \frac{1}{t_{end} - z} \sum_{i=z}^{t_{end}} \frac{y_z - f(z)}{y_z} * 100 \tag{5}$$

To measure the skill of our probabilistic projections we use the Continuous Ranked Probability Score (CRPS) [53] measures "how well the marginal distributions of the forecast represent the ground truth" [54].

For each technology, we compare the performance of the models trained on logistic, bilogistic and mixed data to identify which model most accurately captures empirically observed growth dynamics. We identify the best performing models for onshore wind (bilogistic data), solar PV (mixed data), mobiles (logistic data), and CCGTs (ambiguous), and use these models for subsequent tests and comparisons.

The selection of the best model variant for each technology is based on an evaluation of sMAPE, sMPE and CRPS scores and prediction interval coverage (Supplementary Figures 7, 8). We found that technologies with more complex adoption patterns (like onshore wind, which often shows multi-phase growth) are better captured by the bilogistic model, while technologies with smoother adoption curves (like mobile phones) are better represented by the simpler logistic model. For solar PV, the

mixed model performed best, suggesting its deployment patterns so far show characteristics of both single-phase and multi-phase growth across different countries.

Here, we compare the projections from our main model for each technology to projections made by fitting logistic and exponential curves to global data. These hindcasting tests show that our model using national data out-performs these other models in accurately projecting out-of-sample deployment at the global scale.

Supplementary Note 4: Acceleration scenarios for onshore wind and solar PV

To illustrate the scale of policy effort required to achieve global deployment trajectories in-line with keeping warming below 1.5°C, we develop two sets of stylised counterfactual trajectories. The first, Early Acceleration trajectories assume that the growth of solar PV and onshore wind continues along a logistic curve with additional policy effort accelerating additions by expanding the deployment ceiling. This is achieved by adjusting the logistic curve ceiling (L) to achieve the desired deployment level in 2040 while keeping the intrinsic growth rate (k) fixed. The second, Late Acceleration trajectories assume that the technologies follow our median projection until 2030 and then experience a fresh pulse of accelerating growth modelled using a second logistic curve with another ceiling and a high intrinsic growth rate, essentially producing a bilogistic trajectory. Here, additional policy effort introduces a new growth pulse which alters both, the effective deployment ceiling as well as the rate of growth acceleration.

These counterfactual scenarios represent distinct policy approaches: Early Acceleration reflects immediate, sustained policy intervention that shifts the technology's long-term trajectory by raising it's deployment ceiling and addressing emergent barriers, while Late Acceleration represents a delayed but more intensive intervention that creates a distinct second wave of adoption. By modeling both approaches, we can examine trade-offs between timing and intensity of policy intervention, as well as the feasibility of meeting climate targets through different policy pathways.

These trajectories are compared against Baseline trajectories based on the median probabilistic projections from PROLONG. For onshore wind, this baseline is modelled using a logistic function with $k = 0.16$ and $L = 25\%$. For solar PV, it is modelled using a logistic function with $k = 0.25$ and $L = 18\%$.

For onshore wind, we model the Early Acceleration trajectory using a logistic function with $k = 0.16$ (matching the Baseline) and an increased ceiling $L = 38\%$ (more than $2\times$ higher than the

Baseline). We model the Late Acceleration trajectory using a bilogistic function with an initial pulse with $k_1 = 0.16$ and $L_1 = 25\%$ (same as the Baseline), and a second pulse from 2030 with $k_2 = 0.2$ and $L_2 = 13\%$, which gives an effective L ($L_1 + L_2$) of 38%.

For solar PV, we model the Early Acceleration trajectory using a logistic function with $k = 0.25$ (matching the Baseline) and an increased ceiling $L = 45\%$ (more than 2× higher than the Baseline). We model the Late Acceleration trajectory using a bilogistic function with an initial pulse with $k_1 = 0.25$ and $L_1 = 18\%$ (same as the Baseline), and a second pulse from 2030 with $k_2 = 0.3$ and $L_2 = 27\%$, which gives an effective L ($L_1 + L_2$) of 45%.

For the Early Acceleration trajectories, we:

1. Start with the baseline projection's parameters (k, G, and L)
2. Maintain the growth rate parameter (k) while increasing the ceiling parameter (L)
3. Recalculate the inflection point (t_0) to ensure the curve passes through the current deployment level using the equation:

$$t_0 = \left(\frac{1}{k} \times \log\left(\frac{L}{y_i} - 1\right)\right) + t_i \tag{6}$$

   where y_i is the current deployment level and t_i is the current year
4. Generate the full trajectory using the logistic function:

$$y(t) = \frac{L}{1 + \exp(-k \times (t - t_0))} \tag{7}$$

For Late Acceleration trajectories, we:

1. Take the baseline trajectory up to the second phase start year (2030)
2. Create a second logistic curve with new parameters (k_2, L)
3. Combine the two trajectories using a transition function that shifts from the first to the second curve.

This approach allows us to model the effect of delayed but intensive policy intervention that produces a distinct second wave of technology adoption starting in 2030.

This exercise gives us a set of six global trajectories. For each of these trajectories, we use a linear optimisation approach to distribute the required growth in each year from 2023 to 2050 between 10 regions – East Asia, North America, Europe, South Asia, Asia-Pacific Developed, South-East Asia and Developing Pacific, Africa, Eurasia, Latin America and the Caribbean, and the Middle East. We

introduce constraints on the maximum deployment ceiling for each region for each technology, and on the maximum annual growth rate, and the maximum annual growth acceleration.

The maximum deployment ceilings for each technology in each region are adapted from peak deployment in IPCC AR6 scenarios [55], and have been deliberately set higher than what these scenarios consider achievable in some regions to avoid artificially constraining the optimization (Supplementary Table 7).

The algorithm takes as inputs the current deployment levels ($d_{i,t}$) and growth rates ($r_{i,t}$) for each region i at time t, regional weights (w_i) that sum to unity, and a target global deployment trajectory (D_t^*). It incorporates three key constraints that reflect physical and institutional limitations: maximum allowable deployment in each region ($d_{i,\max}$), maximum annual growth rate ($r_{\max}$), and maximum year-on-year acceleration in growth rates ($a_{\max}$).

For each timestep, the algorithm optimizes regional growth rates through an iterative process that converges when the weighted sum of regional deployments matches the global target within a specified tolerance ε:

$$\left| \sum_i (w_i \times d_{i,t}) - D_t^* \right| < \varepsilon \tag{8}$$

The regional deployment in each year is given by:

$$d_{i,t} = d_{i,t-1} + r_{i,t} \tag{9}$$

subject to the constraints:

$$0 \leq d_{i,t} \leq d_{i,\max} \quad \text{(deployment bounds)} \tag{10}$$

$$0 \leq r_{i,t} \leq r_{\max} \quad \text{(growth rate bounds)} \tag{11}$$

$$-a_{\max} \leq r_{i,t} - r_{i,t-1} \leq a_{\max} \quad \text{(acceleration bounds)} \tag{12}$$

Based on empirical observations of historical technology diffusion rates, we set $r_{\max} = 0.03$ (3 percentage points increase in market share per year) and $a_{\max} = 0.01$ (maximum 1 percentage point change in annual growth rate). These constraints reflect the practical limitations observed in how quickly technologies can be deployed at regional scales.

For each year, the algorithm first calculates the required global increase in deployment ($\Delta D_t^* = D_t^* - D_{t-1}^*$). It then distributes this increase among regions proportionally to their available headroom ($h_{i,t} = d_{i,\max} - d_{i,t}$). The distribution factor ($f_{i,t}$) for each region is calculated as:

$$f_{i,t} = \frac{h_{i,t}}{\sum_j h_{j,t}} \tag{13}$$

where the sum is over all regions j that have not reached their constraints.

To prevent overshooting the global target, the algorithm implements a conservative adjustment mechanism. When $|\Delta D_t^*| < 0.01$, the algorithm applies a damping factor of 0.5 to the growth rate adjustments. If overshooting occurs ($\sum_i (w_i \times d_{i,t}) > D_t$), the algorithm allows growth rates to decline while maintaining $r_{i,t} \geq 0$ and respecting the acceleration constraint $|r_{i,t} - r_{i,t-1}| \leq a_{\max}$.

The algorithm verifies whether the specified global deployment trajectory is achievable under the given constraints by checking that the weighted sum of maximum regional deployments exceeds the global target:

$$\sum_i (w_i \times d_{i,\max}) \geq D_t^* \quad \forall t \tag{14}$$

If this condition is not met, the algorithm terminates, as the regional constraints preclude achievement of the global target.

This approach allows us to generate regional trajectories (Supplementary Figures 11, 14) that help quantify the regional policy effort required to achieve different global growth trajectories while maintaining consistency with empirically-observed limits on how quickly regions can accelerate technology deployment (Supplementary Figures 12, 14).

Supplementary Note 5: The case of offshore wind power

Our analysis in this paper focuses on the growth of onshore wind power and excludes offshore wind primarily because the sample of countries for this technology is not sufficiently large. Although offshore wind power shares many fundamental characteristics with its onshore counterpart, it possesses several distinguishing features that warrant separate consideration. Its comparatively lower technological maturity, heightened technical complexity, generally larger turbine unit sizes, unique investment profiles, and specific geophysical requirements all contribute to potentially different temporal and spatial diffusion patterns.

At the global scale, the growth of offshore wind achieved takeoff in 2011 and as of 2023, accounted for roughly 7% of total global installed wind capacity and less than 0.8% of global electricity generation. Unlike onshore wind, its cross-country diffusion has remained rather limited. As of 2023, it participated in electricity generation in 19 countries, of which we have identified take-off in 11 countries – Sweden (2008), Denmark (2010), the UK (2011), Germany (2015), the Netherlands (2016), Finland & South Korea (both 2018), Belgium (2019), China, Portugal and Vietnam (all 2020).

Of these 11, 8 countries have seen deployment levels exceed 1% of annual electricity generation with Denmark (26.1%), the UK (17.5%), the Netherlands (10.1%), Belgium (9%), Germany (5.1%) and Vietnam (1.3%) showing particularly strong growth (Supplementary Figures 15, 16). Together, these 7 countries are responsible for about 54% of global offshore wind deployment. China leads deployment in absolute terms and contributes to 52% of the global total – this means that over 96% of global offshore wind deployment is concentrated in just 8 countries.

Only Belgium, Denmark, Finland, Germany, South Korea, Sweden and the UK have ≥ 5 datapoints after take-off and are eligible for curve-fitting. As of 2023, only Denmark, Germany and Sweden had exhibited logistic curve maturities over 50%, indicating that growth is still accelerating in most countries. This leaves us with a sample far too small to use with our approach based on using evidence from countries at more mature phases of adoption to inform projections for global growth.

References

[1] Bogdanov, D. *et al.* Low-cost renewable electricity as the key driver of the global energy transition towards sustainability. *Energy* **227**, 120467 (2021).

[2] Way, R., Ives, M. C., Mealy, P. & Farmer, J. D. Empirically grounded technology forecasts and the energy transition. *Joule* **6**, 2057–2082 (2022).

[3] IEA. World energy outlook 2024. Tech. Rep. (2024). URL https://iea.blob.core.windows.net/assets/140a0470-5b90-4922-a0e9-838b3ac6918c/WorldEnergyOutlook2024.pdf.

[4] Luderer, G. *et al.* Impact of declining renewable energy costs on electrification in low-emission scenarios. *Nature Energy* **7**, 32–42 (2022).

[5] Nijsse, F. J. M. M. *et al.* The momentum of the solar energy transition. *Nature Communications* **14**, 6542 (2023).

[6] Griliches, Z. Hybrid corn: An exploration in the economics of technological change. *Econometrica* **25**, 501 (1957).

[7] Rogers, E. M. *Diffusion of Innovations* 3 edn (The Free Press, New York; United States, 1983).

[8] Grubler, A. *The Rise and Fall of Infrastructures Dynamics of Evolution and Technological Change* (Physica-Verlag Heidelberg, Heidelberg; Germany, 1990). URL https://pure.iiasa.ac.at/id/eprint/3351/.

[9] Kazlou, T., Cherp, A. & Jewell, J. Feasible deployment of carbon capture and storage and the requirements of climate targets. *accepted to Nature Climate Change* (2024).

[10] Brutschin, E., Cherp, A. & Jewell, J. Failing the formative phase: The global diffusion of nuclear power is limited by national markets. *Energy Research & Social Science* **80**, 102221 (2021).

[11] Cherp, A., Vinichenko, V., Tosun, J., Gordon, J. A. & Jewell, J. National growth dynamics of wind and solar power compared to the growth required for global climate targets. *Nature Energy* **6**, 742–754 (2021).

[12] Bento, N., Wilson, C. & Anadon, L. D. Time to get ready: Conceptualizing the temporal and spatial dynamics of formative phases for energy technologies. *Energy Policy* **119**, 282–293 (2018). URL https://doi.org/10.1016/j.enpol.2018.04.015.

[13] Bento, N. & Wilson, C. Measuring the duration of formative phases for energy technologies. *Environmental Innovation and Societal Transitions* **21**, 95–112 (2016). URL http://dx.doi.org/10.1016/j.eist.2016.04.004.

[14] Grubler, A., Wilson, C. & Nemet, G. Apples, oranges, and consistent comparisons of the temporal dynamics of energy transitions. *Energy Research & Social Science* **22**, 18–25 (2016).

[15] Bergek, A., Jacobsson, S., Carlsson, B., Lindmark, S. & Rickne, A. Analyzing the functional dynamics of technological innovation systems: A scheme of analysis. *Research Policy* **37**, 407–429 (2008). URL message:%3CCAPfknAVYNJ2orwYcemUZECisCE_AS5mZhGcgJN4i7SiyiTgt-Q@mail.gmail.com%3E.

[16] Jacobsson, S. & Bergek, A. Transforming the energy sector: the evolution of technological systems in renewable energy technology. *Industrial and Corporate Change* **13**, 815–849 (2004).

[17] Markard, J. The next phase of the energy transition and its implications for research and policy. *Nature Energy* **3**, 628–633 (2018). URL http://dx.doi.org/10.1038/s41560-018-0171-7. First, policy support for now mature technologies such as wind and solar is likely to decrease (sic!).

[18] Jacobsson, S. & Lauber, V. The politics and policy of energy system transformation—explaining the german diffusion of renewable energy technology. *Energy Policy* **34**, 256–276 (2006).

[19] Gosens, J., Hedenus, F. & Sandén, B. A. Faster market growth of wind and PV in late adopters due to global experience build-up. *Energy* **131**, 267–278 (2017).

[20] Vinichenko, V., Jewell, J., Jacobsson, J. & Cherp, A. Historical diffusion of nuclear, wind and solar power in different national contexts: implications for climate mitigation pathways. *Environmental Research Letters* **18** (2023).

[21] Malhotra, A. & Schmidt, T. S. Accelerating low-carbon innovation. *Joule* **4**, 2259–2267 (2020).

[22] Wilson, C. *et al.* Granular technologies to accelerate decarbonization. *Science* **368**, 36–39 (2020).

[23] Pierson, P. Increasing returns, path dependence, and the study of politics. *American Political Science Review* **94**, 251–267 (2000).

[24] Jacobsson, S. & Johnson, A. The diffusion of renewable energy technology: an analytical framework and key issues for research. *Energy Policy* **28**, 625–640 (2000).

[25] Ioannidis, R. & Koutsoyiannis, D. A review of land use, visibility and public perception of renewable energy in the context of landscape impact. *Applied Energy* **276**, 115367 (2020).

[26] Frantál, B., Frolova, M. & Liñán-Chacón, J. Conceptualizing the patterns of land use conflicts in wind energy development: Towards a typology and implications for practice. *Energy Research & Social Science* **95**, 102907 (2023).

[27] Mulvaney, D. Identifying the roots of green civil war over utility-scale solar energy projects on public lands across the american southwest. *Journal of Land Use Science* **12**, 493–515 (2017).

[28] Stokes, L. C., Franzblau, E., Lovering, J. R. & Miljanich, C. Prevalence and predictors of wind energy opposition in north america. *Proceedings of the National Academy of Sciences* **120**, e2302313120 (2023).

[29] Ko, I. Rural opposition to landscape change from solar energy: Explaining the diffusion of setback restrictions on solar farms across south korean counties. *Energy Research & Social Science* **99**, 103073 (2023).

[30] Crawford, J., Bessette, D. & Mills, S. B. Rallying the anti-crowd: Organized opposition, democratic deficit, and a potential social gap in large-scale solar energy. *Energy Research & Social Science* **90**, 102597 (2022).

[31] Patterson, J. J. Backlash to climate policy. *Global Environmental Politics* **23**, 68–90 (2023).

[32] Breetz, H., Mildenberger, M. & Stokes, L. The political logics of clean energy transitions. *Business and Politics* **20**, 492–522 (2018).

[33] Stokes, L. C. & Breetz, H. L. Politics in the u.s. energy transition: Case studies of solar, wind, biofuels and electric vehicles policy. *Energy Policy* **113**, 76–86 (2018).

[34] Stokes, L. C. Electoral backlash against climate policy: A natural experiment on retrospective voting and local resistance to public policy. *American Journal of Political Science* **60**, 958–974 (2016).

[35] Gorman, W. *et al.* Grid connection barriers to renewable energy deployment in the united states. *Joule* **9**, 101791 (2025).

[36] Johnston, S., Liu, Y. & Yang, C. An empirical analysis of the interconnection queue (2023).

[37] Heptonstall, P. J. & Gross, R. J. K. A systematic review of the costs and impacts of integrating variable renewables into power grids. *Nature Energy* **6**, 72–83 (2021).

[38] Martinot, E. Grid integration of renewable energy: Flexibility, innovation, and experience. *Annual Review of Environment and Resources* **41**, 223–251 (2016).

[39] Markard, J., Bento, N., Kittner, N. & Nuñez-Jimenez, A. Destined for decline? examining nuclear energy from a technological innovation systems perspective. *Energy Research & Social Science* **67**, 101512 (2020).

[40] Verhulst, P. Recherches mathématiques sur la loi d'accroissement de la population. *Nouveaux Mémoires de l'Académie Royale des Sciences et Belles-Lettres de Bruxelles* **18**, 1–45 (1845).

[41] Wilson, C., Grubler, A., Bauer, N., Krey, V. & Riahi, K. Future capacity growth of energy technologies: are scenarios consistent with historical evidence? *Climatic Change* **118**, 381–395 (2013).

[42] Debecker, A. & Modis, T. Determination of the uncertainties in s-curve logistic fits. *Technological Forecasting and Social Change* **46**, 153–173 (1994). Estimate the level of maturity of the logistic curve which is necessary to reliably estimate K from empirical data.

[43] Odenweller, A., Ueckerdt, F., Nemet, G. F., Jensterle, M. & Luderer, G. Probabilistic feasibility space of scaling up green hydrogen supply. *Nature Energy* **7**, 854–865 (2022).

[44] Grubb, M., Drummond, P. & Hughes, N. *The shape and pace of change in the electricity transition* (We Mean Business Coalition, Washington, DC; USA, 2020). URL https://www.wemeanbusinesscoalition.org/wp-content/uploads/2020/10/Shape-and-Pace-of-Change-in-the-Electricity-Transition-1.pdf.

[45] Dixon, R. Hybrid corn revisited. *Econometrica* **48**, 1451 (1980).

[46] Gompertz, B. On the nature of the function expressive of the law of human mortality and on a new mode of determining the value of life contingencies. *Phil.Trans.Roy.Soc.* **123**, 513–585 (1825).

[47] Haegel, N. M. *et al.* Photovoltaics at multi-terawatt scale: Waiting is not an option. *Science* **380**, 39–42 (2023).

[48] Victoria, M. *et al.* Solar photovoltaics is ready to power a sustainable future. *Joule* **5**, 1041–1056 (2021).

[49] Creutzig, F. *et al.* The underestimated potential of solar energy to mitigate climate change. *Nature Energy* **2**, 17140 (2017).

[50] Zielonka, N. & Trutnevyte, E. Probabilities of reaching required diffusion of granular energy technologies in european countries. *iScience* **28**, 111825 (2025).

[51] Zielonka, N., Wen, X. & Trutnevyte, E. Probabilistic projections of granular energy technology diffusion at subnational level. *PNAS Nexus* **2**, pgad321 (2023).

[52] Wen, X., Jaxa-Rozen, M. & Trutnevyte, E. Accuracy indicators for evaluating retrospective performance of energy system models (2022).

[53] Gneiting, T. & Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* **102**, 359–378 (2007).

[54] Price, I. *et al.* Probabilistic weather forecasting with machine learning. *Nature* **637**, 84–90 (2025).

[55] Byers, E. *et al.* AR6 scenarios database (2022). URL https://doi.org/10.5281/zenodo.5886912.