Will national renewable costs continue declining?

C. Lennart Baumgärtner[1,2,3] and J. Doyne Farmer[1,2,4,5]

[1]Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, OX1 3UQ, UK
[2]Smith School of Enterprise and the Environment, University of Oxford, Oxford, OX1 3QY, UK
[3]Vallorii Ltd, 450-500 Silbury Boulevard, Milton Keynes MK9 2AD, UK
[4]Macrocosm Inc, 235 E. 4th St., Brooklyn NY 11218
[5]Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM 87501

March 10, 2026

Supplementary Information

S1. Notation overview

Variable	Name	Description	Unit
$\mathcal{L}$	Levelized Cost of Electricity (LCOE)	Average annual cost to produce 1 kWh of electricity	USD/kWh
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 I_0, 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 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
γ	O&M cost over initial investment cost	Simplified notation for $\gamma = O/I$	$\%/year$
$\mathcal{M}$	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}
ι	Inverter loading ratio	Loading ratio or capacity of solar PV inverter	$kW_{peak}/kW_{nameplate}$
$\mathcal{A}$	Set of LCOE components	This set includes LCOE components $A \in \mathcal{A} = \{I, \mathcal{C}, O, \mathcal{Y} + \gamma\}$	n.a.
$\mathcal{E}$	Average LCOE error	Multiplicative error term between the average national LCOE and the LCOE of the national averages	n.a.

Table S1: **Overveiew of notation used throughout this article.**

S2. Description of national data collection and cleaning

S2.1 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[1], IEA[2], and IRENA[3]. 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.

S2.2 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.

S2.2.1 *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[4].

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.[5] 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.

This means that we are taking the WACC as an exogenous variable rather than making explicit forecasts for it. Given that the WACC is mostly driven by the wider macroeconomic environment, forecasting renewable-specific WACCs is out of scope for this work.

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 S3.7, 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. IRENA applies strong assumptions about the O&M costs in their dataset. These are published for data post 2010. For data prior to 2010 we can calculate the implied O&M costs based on other data they provide, as well as the LCOE formula they use. IRENA's O&M cost assumptions before 2010 differ from other data sources and our findings in Section S3.8. We replace their assumptions with our findings to avoid artificially increasing the noise in our forecasts.

We thus do not validate our forecasts for the O&M costs prior to 2010. After 2010, our forecasts are based on a simple relationship with the total investment costs described later. This simplified approach reflects the relatively small contribution of O&M costs to solar LCOE, poor data quality, high degree of cross-sectional correlation, and high degree of correlation with the total investment cost. If more data becomes available on O&M costs in the future, this would be a promising avenue of research to further improve our methods.

S2.2.2 *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 different discount rate to the electricity generation to represent differences in future social preferences.

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) and governmental (all-in) Feed-In Tariffs (FITs). 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.[6] analyze 347 data points for national solar BOS costs (although they differentiate commercial and residential solar, almost doubling their dataset), Kothari et al.[7] use 112 data points for national BOS costs, Riva et al.[8] analyze 52 data points for wind O&M costs.

S2.2.3 *Different LCOE formulae*

There exist two commonly used formula for the LCOE $\mathcal{L}_{ijt}$. In particular, some research applies a discount rate to future electricity generated by a renewable asset:

$$\mathcal{L}_{ijt} \equiv \frac{I_{ijt} + \sum_{\tau=1}^{T} \frac{O_{ijt}}{(1+\mathcal{D}_{ijt})^{\tau}}}{\sum_{\tau=1}^{T} \frac{\mathcal{Y}_{ijt}}{(1+\mathcal{D}_{ijt})^{\tau}}}, \tag{S1}$$

where T is the asset lifetime, O_{ijt} is the operating and maintenance cost in year τ, $\mathcal{Y}_{ijt}$ is the electricity generated in year τ, and $\mathcal{D}_{ijt}$ is the discount rate set equal to the weighted average cost of capital (WACC) discussed later. All terms are set relative to the capacity of the asset. We fix O_{ijt} and $\mathcal{Y}_{ijt}$ for all years τ, so that we can write

$$\mathcal{L}_{ijt} \equiv \frac{I_{ijt} * \sum_{\tau=1}^{T} \frac{1}{(1+\mathcal{D}_{ijt})^{\tau}} + O_{ijt}}{\mathcal{Y}_{ijt}}, \tag{S2}$$

where $\sum_{\tau=1}^{T} \frac{1}{(1+\mathcal{D}_{ijt})^{\tau}}$ is exactly the capital recovery factor $\mathcal{C}_{ijt}$, proving the equivalence of both formulae.

S2.3 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 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.

S2.4 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 (PVPSP).[9] 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.[10]

Turbine data is based on IRENA,[4] 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.[11,12]

S2.5 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.[6] 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

$$\mathcal{I}_{jt} = \mathcal{M}_{jt-\tau} + \mathcal{J}_{jt}. \tag{S3}$$

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.[13,14] 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.[15] 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 τ.

S2.6 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

$$\mathcal{Y}_{ijt} = \frac{Y_{ijt}}{X_{ijt}}, \tag{S4}$$

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. In Section S3.7 we compare the out-of-sample accuracy of both methods. For wind, we find that the method based on the LCOE is less accurate. 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.

S2.7 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.[16] 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.[4,8,16–18]

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 total investment costs in German onshore wind in 2012,[8,19] while Raupach-Sumiya et al. calculate almost 11%.[20] In section S3.8 we combine both sources and make our own estimates of O&M costs relative to the total investment costs. Our estimates for solar are around 1%, for wind around 3%.

S2.8 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 I_{ijt} that covers the expected annual returns of equity- and debt-investors. We calculate the CRF from the WACC as[21]

$$\mathcal{C}_{ijt} = \frac{1}{\sum_{t=1}^{T_i} \frac{1}{(1+\mathcal{D}_{ijt})^t}} = \frac{(1+\mathcal{D}_{ijt})^{T_i}\mathcal{D}_{ijt}}{(1+\mathcal{D}_{ijt})^{T_i}-1}. \tag{S5}$$

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:[22]

$$\mathcal{D}_{jt}^{k} = r_{jt} + p_k. \tag{S6}$$

We can infer the risk-free rate from macroeconomic data. Specifically, we use long-term (10 and 7-year) government bond yields from Bloomberg[23] and Reuters Eikon.[24] For countries without data, we use proxies from IMF.[25] 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.[21] 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.[26] 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[27] 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.[22] 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.[28] For South America and the Middle East, where data is particularly limited, we use the theoretically derived data in Ondraczek et al.[29]

There are two important assumptions underlying equations (S5) and (S6). 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, the asset lifetime may be longer than the financing period, i.e., debt may be repaid before the asset ceases operation. Our CRF figures can be re-scaled to take this into account if our forecasts are to be applied to such assets. 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[30] shows that this is appropriate for current and historical projects. When assets are not funded through project finance, the WACC needs to be adjusted depending on the credit rating of the corporate funding the asset.

S3. LCOE model development

The model development consists of two key steps:

1. Decompose the asset-level LCOE $\mathcal{L}_{ijt}$ into its components $A_{ijt} \in \{I_{ijt}, \mathcal{C}_{ijt}, O_{ijt}, \mathcal{Y}_{ijt}\}$.
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 S3.2 and S3.3 describe the PV module and wind turbine model development, respectively. Sections S3.4 and S3.5 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 S3.6. Sections S3.8 provides additional details on the operating and maintenance cost model.

S3.1 LCOE decomposition and aggregation

Summary: *We decompose the non-linear LCOE into its components using a linear approximation. 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

$$\mathcal{L}_{ijt} \equiv \frac{I_{ijt} * \mathcal{C}_{ijt} + O_{ijt}}{\mathcal{Y}_{ijt}} \quad \text{(USD / kWh)}, \tag{S7}$$

for asset i in country j and year t. We are ultimately interested in the capacity-weighted average LCOE $\mathcal{L}_{jt}$. We average over assets i, for each country j and year t, specific to the technology. Each component model forecasts future values of the capacity-weighted average component $(I_{jt}, \mathcal{C}_{jt}, O_{jt}, \mathcal{Y}_{jt})$. For simplicity, we drop the index t for the remainder of this section. Unfortunately, calculating $\mathcal{L}_j$ from the capacity-weighted average components is not straightforward due to the non-linear nature of Equation (S7). 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 (S7) 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.[31] to confirm that we have not grossly over-fitted our data.

S3.1.1 *Linear LCOE approximation*

In this section, we asses the relationship between the weighted average LCOE $\mathcal{L}_j$ and the weighted average components $\{I_j, \mathcal{C}_j, O_j, \mathcal{Y}_j\}$, particularly:

$$\mathcal{L}_j \stackrel{?}{=} \mathcal{L}(I_j, \mathcal{C}_j, O_j, \mathcal{Y}_j). \tag{S8}$$

Here, we develop a theoretical model of this relationship that we later test empirically.

We first consider the LCOE of a single asset i, dropping the subscript t for convenience. As discussed in the main text and section S3.8, we model the operating and maintenance costs as a fixed share γ of the total investment costs, i.e., $O_{ij} = \gamma * I_{ij}$. This means we can re-write the asset-level LCOE (S7) as a linear equation:

$$\log \mathcal{L}_{ij} = \log I_{ij} + \log (\mathcal{C}_{ij} + \gamma) - \log \mathcal{Y}_{ij} = \sum_{A \in \mathcal{A}} \log A, \tag{S9}$$

for components $A \in \{I, \mathcal{Y}^{-1}, O, \mathcal{C} + \gamma\} =: \mathcal{A}$. 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

$$\overline{\log X} \approx \log \bar{X} + \varepsilon_X, \tag{S10}$$

where the first-order error term ε_X depends on the variance in the random variable X. In particular, if X has a low variance, ε_X is small, but if the variance is high, ε_X is large.

Using this identity, we can write Equation (S9) as

$$\log \mathcal{L}_j \approx \log I_j + \log(\mathcal{C}_j + \gamma) - \log \mathcal{Y}_j + \varepsilon_I + \varepsilon_{\mathcal{C}} - \varepsilon_{\mathcal{Y}} - \varepsilon_{\mathcal{L}} \tag{S11}$$

$$\approx \log I_j + \log(\mathcal{C}_j + \gamma) - \log \mathcal{Y}_j + \log \mathcal{E}_j, \tag{S12}$$

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.

S3.1.2 *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 $\log I_j \mathcal{Y}_j^{-1}(\mathcal{C}_j + \gamma)$ from Equation (S12). We first perform a regression analysis on Equation (S12). This tells us how much of the variance of the log-LCOE is explained by the log of 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,[32] we find that most extreme values in $\log \mathcal{E}_j$ are caused by outliers. Once we remove the outliers, which we believe to be spurious, the distribution can be well-approximated as a normal distribution.

Regression analysis We consider the model

$$\log \mathcal{L}_j = \beta_I \log I_j + \beta_{\mathcal{C}} \log(\mathcal{C}_j + \gamma) - \beta_{\mathcal{Y}} \log \mathcal{Y}_j + \epsilon_j. \tag{S13}$$

Table S2 shows the results of the regression. Equation (S13) 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 I and capacity factor $\mathcal{Y}$, we can already explain 97% of the variance.

Figure S1 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 centered 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

	Solar					Wind				
$\log I$	-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 S2: Regression results for LCOE against its components, according to Equation (S13). We omit the indication of statistical significance since all regressors are highly significant above 99%.

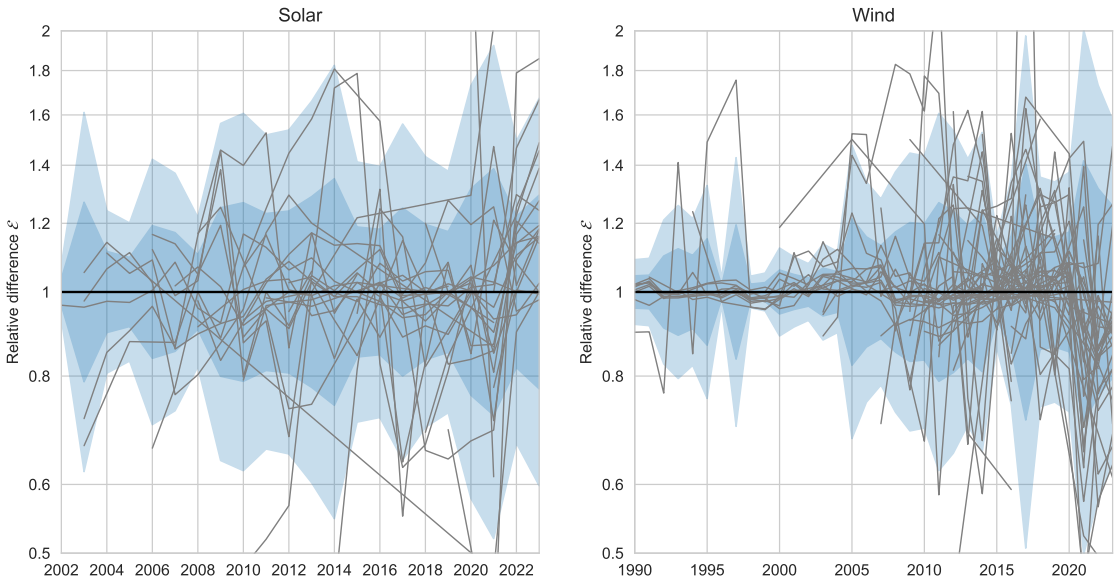

Figure S1: **Natural Logarithm of the relative difference between the average LCOE $\mathbb{E}_\omega \mathcal{L}$ and the LCOE of the averages $\hat{I}\hat{\mathcal{Y}}(\hat{\mathcal{C}} + \hat{\gamma})$ over time.** Grey lines show $\mathcal{E}$ for a specific country over time. Blue shading indicates 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.

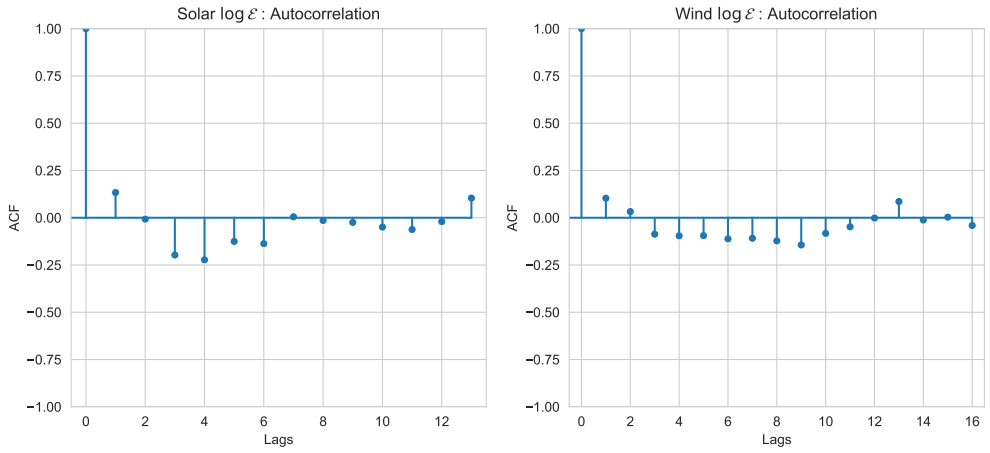

Figure S2: **Average national autocorrelation function** $\log \mathcal{E}$**.** We take the average of the ACF of $\log \mathcal{E}$ for individual countries for solar PV **(left)** and wind **(right)**. Neither technology shows significant autocorrelation for the first lags.

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 time series $(X_t)_{t\in\mathbb{N}}$ and lag τ 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 S2, 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 S3, 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 tails fits this distribution well, as indicated by the orange line in Figure S3. 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 $\log I_{ijt}$ in the US utility-scale solar sector using an asset-level dataset from Berkley Lab.[32] The Berkley dataset consolidates the investment costs for over 2.4 million solar photovoltaic assets in the United States between 2000 and 2023. It covers around 5% of total U.S. solar capacity in the US, although coverage varies substantially by year.

Figure S4 (right panel) 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 (red dots), with a log-difference of up to 20%. The reason

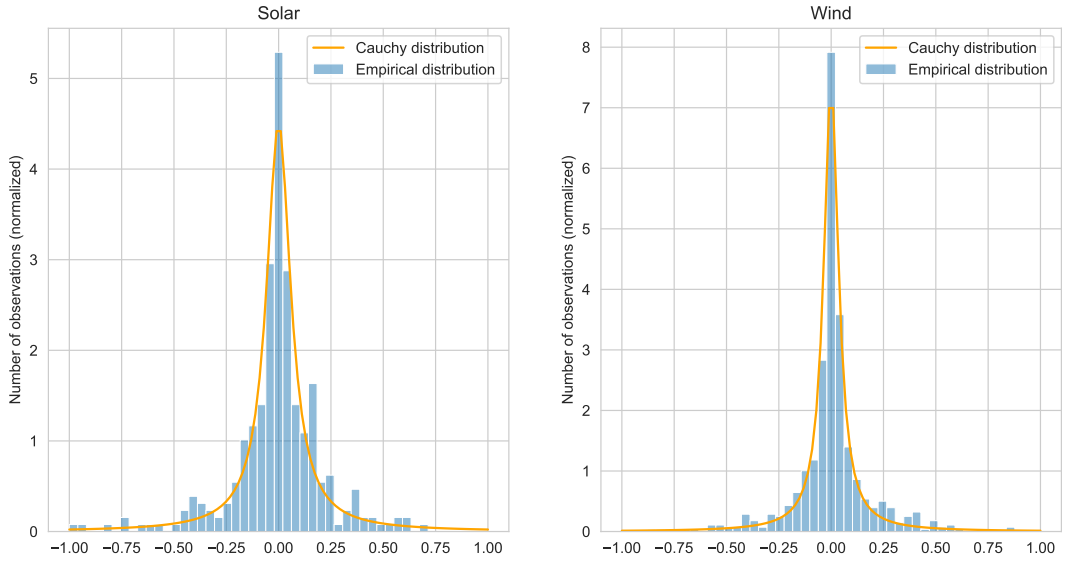

Figure S3: **Empirical distribution of** $\log \mathcal{E}$**.** We calculate $\log \mathcal{E}$ for solar **(left)** and wind **(right)** and look at the resulting distribution. In both cases, we observe a large number of very positive or negative values of $\log \mathcal{E}$, indicating a heavy-tailed distribution.

for this is individual renewable assets with large capacity and particularly low costs. The left panel of Figure S4 shows the distribution of the total investment cost across the Barkley Lab sample, excluding datapoints with costs above USD 3 million. There are a number of utility-scale assets with costs below USD 400 thousand that are visible outliers to the sample distribution. If we remove utility-scale assets with unreasonably low investment costs below USD 400 thousand, average and log-average closely align (blue dots on the right panel).

Including these outliers in our forecasts inflates the errors of our LCOE forecast. We, therefore, exclude extreme $\log \mathcal{E}_{jt}$ values when estimating the size of this error term. Based on the previous total investment cost analysis for the US, we exclude $\log \mathcal{E}_{jt}$ differences exceeding 20%. This is approximately equal to excluding the 10th and 90th percentile. The remaining errors are forecast with a normal distribution, based on a standard deviation estimate of roughly 9% for solar and 7% for wind. This is a rather simple analysis that requires further asset-level justification when such data becomes available.

S3.1.3 *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.[33] 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.[34] 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 are not statistically significant. This insignificance stems from the fact that we do not have sufficient data for the out-of-sample test results to be significant. This is because we can only use a subset of our data for this test. As we will discuss in section S3.6, we still expect the disaggregate model to outperform the aggregate when training it on our full dataset if Wright's or Moore's law continue to hold for the LCOE components.

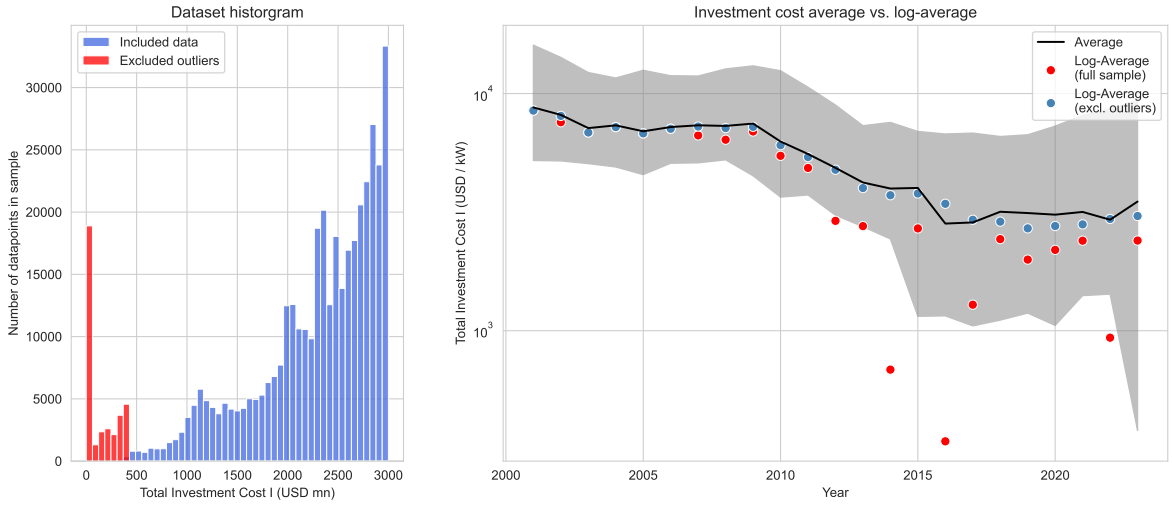

Figure S4: **Average and log-average of the total investment cost I in the US. (Left panel)** The distribution of total investment cost from the Berkley Lab dataset.[32] We exclude datapoints with costs above USD 3mn for better visualization. There are a number of datapoints with costs below USD 400 tsd that are outliers with respect to the full data distribution. **(Right panel)** The black line shows the average 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 $\log I_{ijt}$ over time. For most years, it falls close to I_{jt}, within the 95th quantile. There are, however, multiple outliers with an order of magnitude lower log-average costs. The blue dots show the average of $\log I_{ijt}$ if we exclude the outliers with costs below USD 400 tsd.

S3.2 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.[31] to forecast future global average module costs, conditional on global solar deployment.*

Most countries have had very similar module costs in the past.[35] This is expected with such a highly consolidated market; China has around 80% global market share in module manufacturing.[9] One exception to this is Japan. Protectionism in the Japanese market resulted in the decoupling of local prices from the global average after 2010.[20] We thus exclude Japan from the module cost forecast. 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.[31] 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, includingn solar modules. We therefore apply this model to project future solar module costs:

$$\text{Module model:}\quad \log \mathcal{M}(t) = \log \mathcal{M}(t-1) - \hat{\omega} \log \frac{Z(t)}{Z(t-1)} + u_t + \rho u_{t-1}, \tag{S14}$$

where $Z(t)$ is the global solar capacity until time t, ω 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 (S14) has been extensively studied on over 50 technologies, as well as applied to broader energy system scenarios.[36]

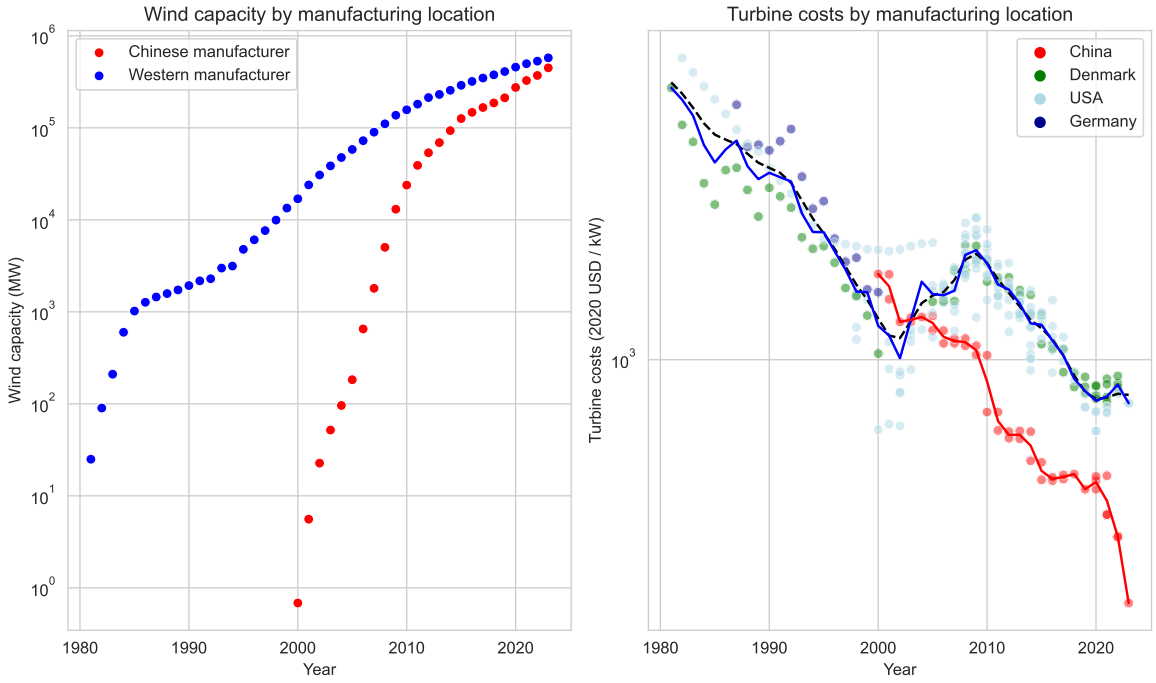

Figure S5: **Wind turbine costs vs cumulative deployment and year.** The left panel shows the increase in wind turbine capacity manufactured by Chinese and Western companies between 1982 and 2021. The right panel shows the decrease in the respective turbine costs over the same period. The red and blue 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.

S3.3 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 is more accurate than a global model. We show that the stochastic Wright's law model of Lafond et al.,*[31] *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 S5, Western-manufactured (Europe and the US) and Chinese-manufactured turbines used to have similar costs in the early 2000s. However, 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.[11,12,37,38]

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 must select one of these four alternatives to make forecasts. Here, we choose a global experience curve model in line with Lafond et al,[31] due to its simplicity. A global model simplifies the IAM forecast comparison performed later. 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.

To validate our model on regional turbine costs, 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 (S14) for each market j (either Western or Chinese turbines):

$$\text{Turbine model: } \log \mathcal{M}_j(t) = \log \mathcal{M}_j(t-1) - \omega \log \frac{Z(t)}{Z(t-1)} + \eta_t^j + \rho \eta_{t-1}^j, \tag{S15}$$

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 ω 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 $\rho = 0.19$.[31]

We estimate the model parameters $\hat{\omega}$ and $\hat{\sigma}_\eta$ with the same method as Lafond et al.[31] 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 ω nor the exogenous shock scale σ_η^2 are specific to the turbines being manufactured in China or the West. Both manufacturing markets share universal parameters since they follow the same learning dynamics. Secondly, the bias-variance trade off means that fewer model parameters are preferred in cases with little data. Smart, Lafond & Farmer[39] show this explicitly for Wright's law. With short time-series 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 τ. We can now compare the theoretical uncertainty distribution to the empirically observed residuals for each value of t, τ, and market j.

To compare these probabilistic forecasts to the true data, we compute the cumulative *probability integral transform* (PIT) of our forecasts.[40,41] 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 S6 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. This difference, however, is not statistically significant.

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 S6. 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 the graphical analysis. We see that our gray 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.

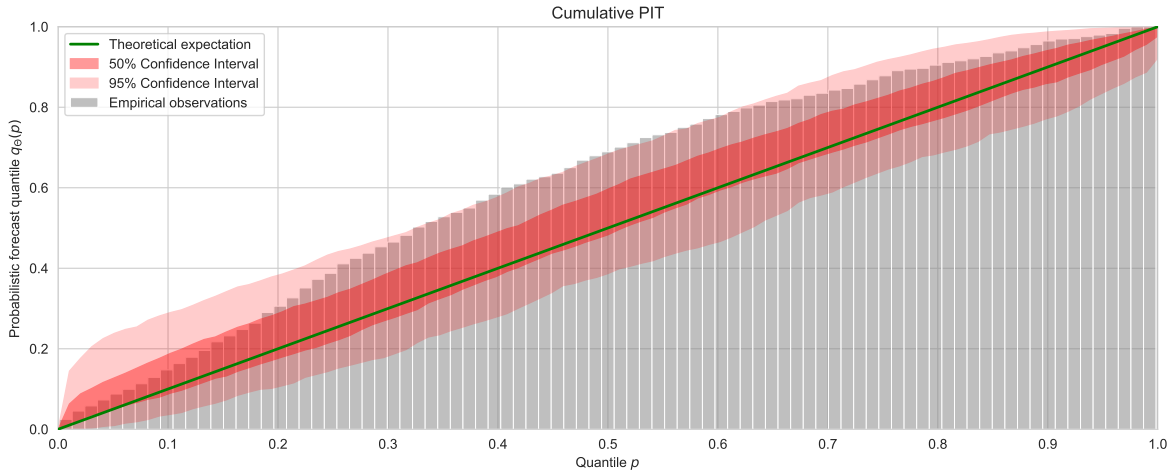

Figure S6: **Out-of-sample validation of the turbine cost forecast error.** We estimate model parameters on 5-year moving windows of our data. We make stochastic forecasts for the future of this window with a Monte Carlo approach. The empirical cumulative PIT $q_\Theta(p)$ is shown in gray. The green 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.

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.

S3.4 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.

S3.4.1 *In-sample analysis*

We first perform a descriptive analysis of the solar BOS costs data. National 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 et al.[42,43] We estimate the statistic

$$CD = \sqrt{\frac{1}{N(N-1)}} \left(\sum_{j=1}^{N-1} \sum_{j'=j+1}^{N} \sqrt{T_{jj'}} \hat{\rho}_{jj'} \right), \tag{S16}$$

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 Metrology.[44] The ESDOF shows that out of the 42 national BOS time series (59 for wind), only 5 (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 calculate 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[44]

$$N^* = \frac{\left(\sum \lambda_i\right)^2}{\sum \lambda_i^2}. \tag{S17}$$

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

$$\log \overline{\mathcal{J}}_j(t) = \log \mathcal{J}_j(t) - \log \mathcal{J}(t), \tag{S18}$$

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.

For solar, the resulting time-series $\log \overline{\mathcal{J}}_j(t)$ are no longer cross-sectionally correlated, as confirmed by another cross-sectional dependence test with p-value below 1%. The ESDOF has increased from 5 to 15. This is still significantly below the number of countries in the sample due to the sparse nature of our dataset. Out of the 42 countries in our dataset, only 22 have more then 10 years of available data. For wind, the resulting time-series $\log \overline{\mathcal{J}}_j(t)$ has 22 ESDOF since more data is available. The cross-sectional dependence test still shows that the time-series are correlated, althouth this is likely from the autocorrelation in $\log \overline{\mathcal{J}}_j(t)$.

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 $x_{t-1}, ..., x_{t-h+1}$. 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 measurement error for the (P)ACF. The resulting average functions are shown in Figure S7. Both ACF and PACF show one significant lag. This indicates that our national time series are significantly autocorrelated.

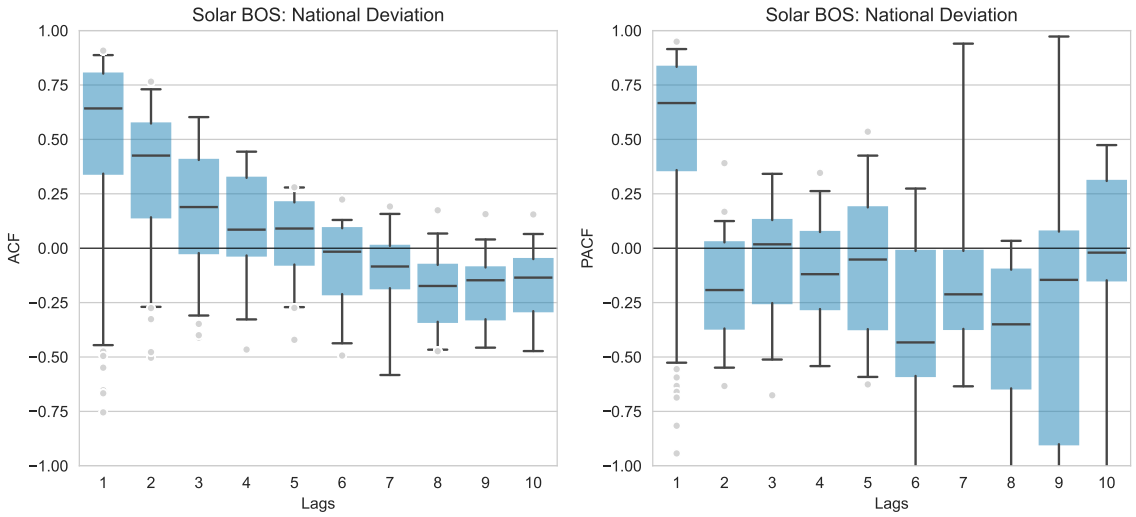

Figure S7: **Cross-sectional weighted average ACF and PACF for the national solar BOS costs.** We calculate national ACF and PACF functions and take their cross-sectional median and confidence intervals. The weights are based on the length of each national time series. **(left panel)** The cross-sectional 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.

S3.4.2 *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.

$$\log \mathcal{J}(t) = \log \mathcal{J}(t-1) - \omega \log \frac{Z(t)}{Z(t-1)} + u_t + \rho u_{t-1}, \tag{S19}$$

$$\log \mathcal{J}_j(t) - \log \mathcal{J}(t) =: \log \overline{\mathcal{J}}_j(t) = \varphi \log \overline{\mathcal{J}}_j(t-1) + \eta_{jt}. \tag{S20}$$

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 (S19). The national deviations from the global average $\log \overline{\mathcal{J}}_j$ are calculated by subtracting the global average from the national time series. The resulting time series is modelled with an AR(1) process, Equation (S20). We estimate φ 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 & Lafond, where[45]

$$\hat{K}_\eta^2 = \hat{\sigma}_\eta^2 / (1 - \varphi^2). \tag{S21}$$

As shown in Figure S8, the cross-sectional variance $\hat{K}_\eta^2$ is decreasing over time. Note that this variance is that of $\log \overline{\mathcal{J}}_j$, 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.

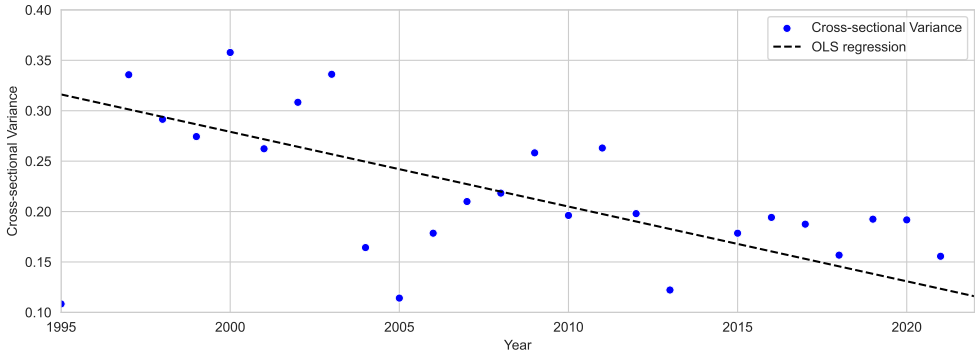

Figure S8: **Cross-sectional variance of the logarithm of the solar BOS over time. (blue dots)** For each year, we compute the cross-sectional variance in the logarithm of national solar BOS costs, $\log \overline{\mathcal{J}}_j(t)$. The black line shows a linear regression of the resulting data. It shows a statistically significant decline.

We focus our model validation on the national deviations from the global average, i.e., $\log \overline{\mathcal{J}}_j$ in Equation (S20) since the global model has been validated by Lafond et al.[31] Equivalent to the turbine model described in Appendix S3.3, we validate this model using backtesting. For each point in time, we make stochastic forecasts for future values of $\log \overline{\mathcal{J}}_j$, 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 S9 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.

S3.4.3 *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,[6] 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 S3 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 (S19)-(S20). 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[39] 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:

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 ω is estimated using a linear regression on the global average.[36]	$\log \mathcal{J}_j(t) = \log \mathcal{J}(t-1) - \omega \log \frac{Z(t)}{Z(t-1)}$	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 ω 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 θ 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 ω_i for each country based on a regression model.[6]	$\log \mathcal{J}_j(t) = \log \mathcal{J}_j(t-1) - \omega_j \log \frac{Z_j(t)}{Z_j(t-1)}$	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 S3: **Overview of compared forecast models for solar BOS costs.** We compare 4 different forecast models forecasting future solar costs in different countries. θ and ω are Moore's and Wright's law parameters, respectively, independent of the market. ω_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

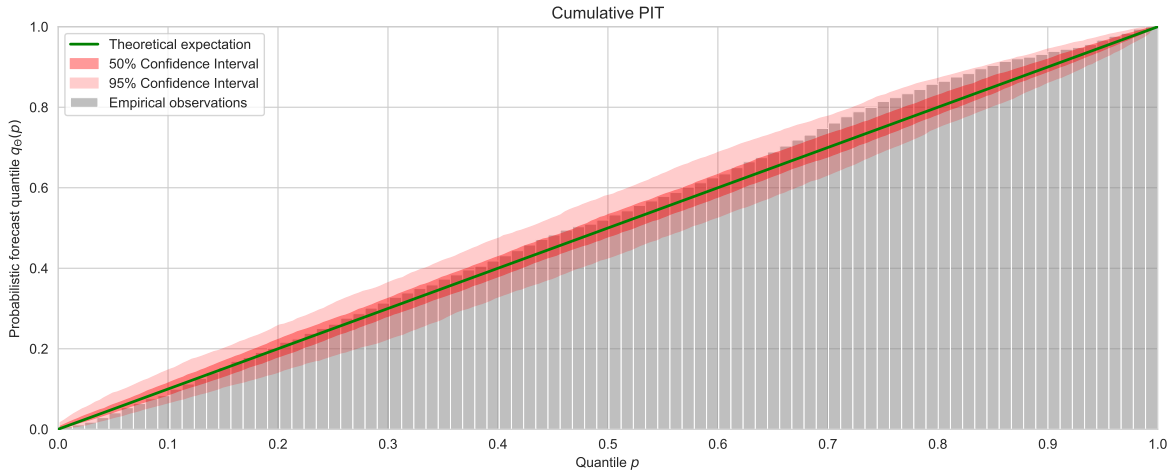

Figure S9: **Out-of-sample validation of the solar BOS cost forecast $\log \overline{\mathcal{J}}_j$, relative to the true global average cost.** We consider only the national solar BOS cost relative to the global average BOS cost $\log \overline{\mathcal{J}}_j$. We make stochastic forecasts for the 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 gray. The green 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.

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*;[46] 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 model l and forecast horizon τ is defined as

$$RMSE(\tau) = \sqrt{\frac{1}{n}\sum_{t,j}\left(\log \hat{\mathcal{J}}_j^l(t+\tau) - \log \mathcal{J}_j^l(t+\tau)\right)^2}, \tag{S22}$$

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 S10. 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.[39,47] 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 & Farmer[39] that compare geometric random walk forecast models with and without parameter heterogeneity. Under the assumption that the national learning parameters ω_i follow a gamma distribution with parameters α_ω and β_ω, 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

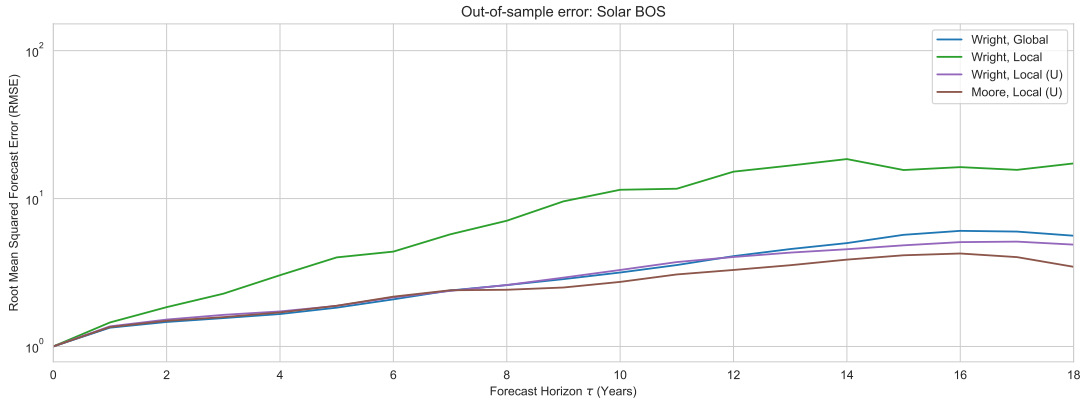

Figure S10: **Out-of-sample residuals for different forecasting models of solar BOS.** BOS costs are forecast under different models described in Table S3. We estimate the model parameters on 5-year moving average training data and compare the forecasts with the actual data. The x-axis shows the forecast 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.

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 ω_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 ω_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 S11 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 Empirically, as we have seen in Figure S10, 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.

There are, however, two additional 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 exercise, we show that the uncertainty associated with the *Wright, Local (U)* model is larger than what we observe in the data. In other words, the uncertainty bands of the *Wright, Local (U)* model are unnecessarily large. Secondly, the *Wright, Local (U)* implies an increase in the cross-sectional variance over time since independent random walks naturally disperse. This assumption is not supported by the data. We investigate both of these facts in turn.

We first need to define the of out-of-sample residuals under the different models. The model used to forecast residuals in the *Wright, Global* model is described in the main text and Section S3.4.2. We

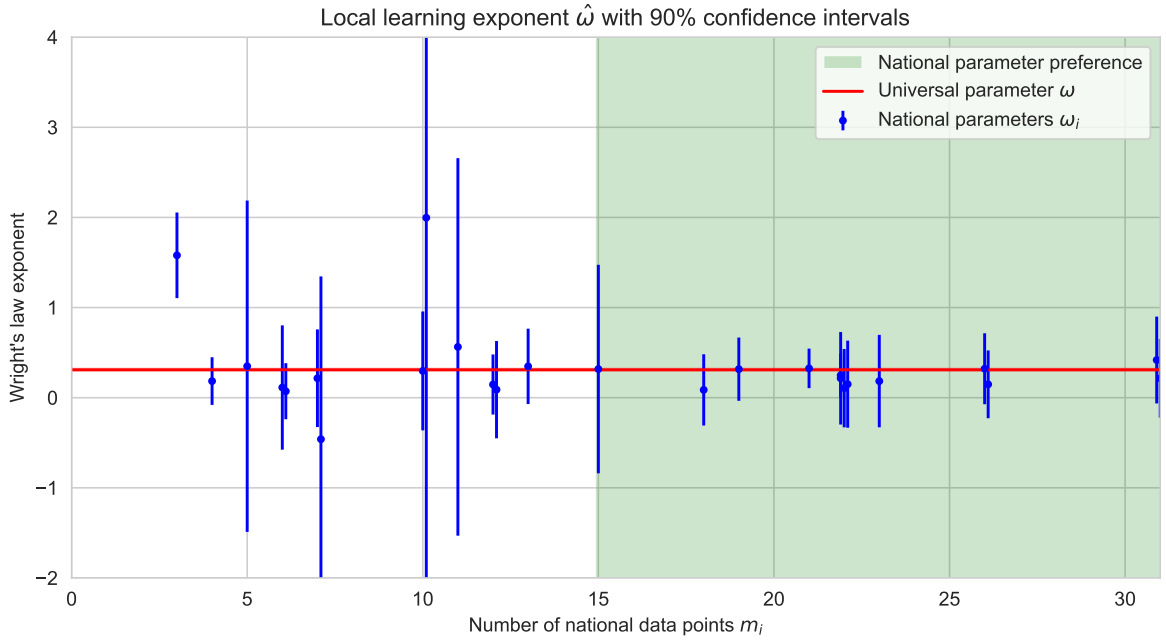

Figure S11: **National Wright's law parameters ω_j against the number of available data points.** We show the estimated national Wright's law exponents $\hat{\omega}_j$ depending on the number of national data points available (m_j). Blue dots indicate the parameters with the associated 90% confidence 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 $\hat{\omega}_j$.

apply the random walk model of Lafond et al.[31] 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.[31] on the national BOS costs:

$$\log \mathcal{J}_j(t) = \log \mathcal{J}_j(t-1) - \hat{\omega} \log \frac{Z_j(t)}{Z_j(t-1)} + \eta_{jt} + \rho \eta_{jt-1}, \tag{S23}$$

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 ω. $\hat{\omega}$ 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

$$CRPS(F, y) = \int_{\mathbb{R}} \left(F(x) - \mathbb{I}_{(x \leq y)} \right)^2 dx, \tag{S24}$$

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 for the CRPS. In our case, we apply the so-called "energy formulation". For each data point y that we make stochastic forecasts for, we write

$$\widehat{CRPS}(\{x_i\}_{i=1,\dots,M}, y) = \frac{1}{M} \sum_{i=1}^{M} |x_i - y| - \frac{1}{2M^2} \sum_{i,j=1}^{M} |x_i - x_j|, \tag{S25}$$

where $\{x_i\}_{i=1,\dots,M}$ is the stochastic forecast ensemble. y is the empirical out-of-sample data we.[48] We choose $M = 500$ in the estimation below. It is non-trivial to show that Equation (S24) and (S25) are equivalent. The first term in (S25) is the mean absolute error (MAE). 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 MAE and low forecast width. A narrow yet inaccurate forecast may have low width yet high MAE. A wider forecast will have a large width and high MAE, 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 (S25) is biased, it outperforms unbiased alternatives for large values of M and few observations y.[48]

As for the RMSE, we compare the average CRPS of our two models for different forecast horizons τ. As explained earlier, we focus on the forecast residual relative to the global average cost, $\log \overline{\mathcal{J}}_j$ defined in Equation (S18), since this significantly increases the statistical significance of our comparison.

The results of this analysis are shown in Figure S12. 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 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

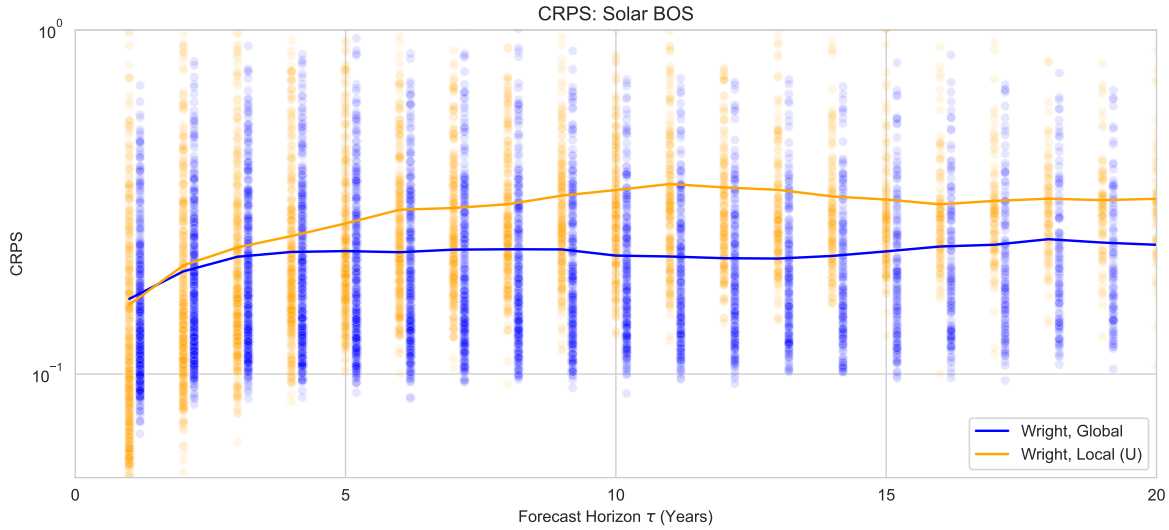

Figure S12: **Out-of-sample comparison of the solar BOS cost forecast errors under the *Wright, Local (U)* and *Wright, Global (U)* model relative to the global average solar BOS costs, in terms of CRPS.** We consider only the national solar BOS cost relative to the global average BOS cost. We estimate 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.

noise is independent for each country, resulting in further divergence of national costs. In reality, we have seen in S8 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.

S3.5 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 mean as a mean-reverting process.*

S3.5.1 *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

$$\log \mathcal{J}_j(t) = \alpha_j + \omega * \log Z_j(t) + \tilde{\omega} * \log Z(t). \tag{S26}$$

α_j is a country fixed effect, ω 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 S4. For solar, we see highly significant global and local learning. This is expected given the large cost decline and take up of solar

	Solar PV			Wind		
Local learning ω	-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 S4: Regression results for BOS costs, according to Equation (S26). For Solar, all regression parameters are highly significant. For wind, the local learning parameter is significant below 10% if we include only local learning.

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.[49] 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 inflation-adjusted prices have remained approximately constant for more than a century despite improvements in extraction technologies.[50] 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 et al.,[42] see Equation (S16). 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 S13. 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 (S18). We then calculate the ACF and PACF for $\log \overline{\mathcal{J}}_j(t)$ 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 panels of figure S14. We find one significant lag in both the ACF or PACF. This means that we can approximate the $\log \overline{\mathcal{J}}_j(t)$ 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

$$\mathcal{J}_j(t) = I_j(t) - \omega_{t-1}\mathcal{M}_j(t-1) - \omega_{t-2}\mathcal{M}_j(t-2), \tag{S27}$$

where $I(t)$ 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 ω_t or, equivalently, different lags.

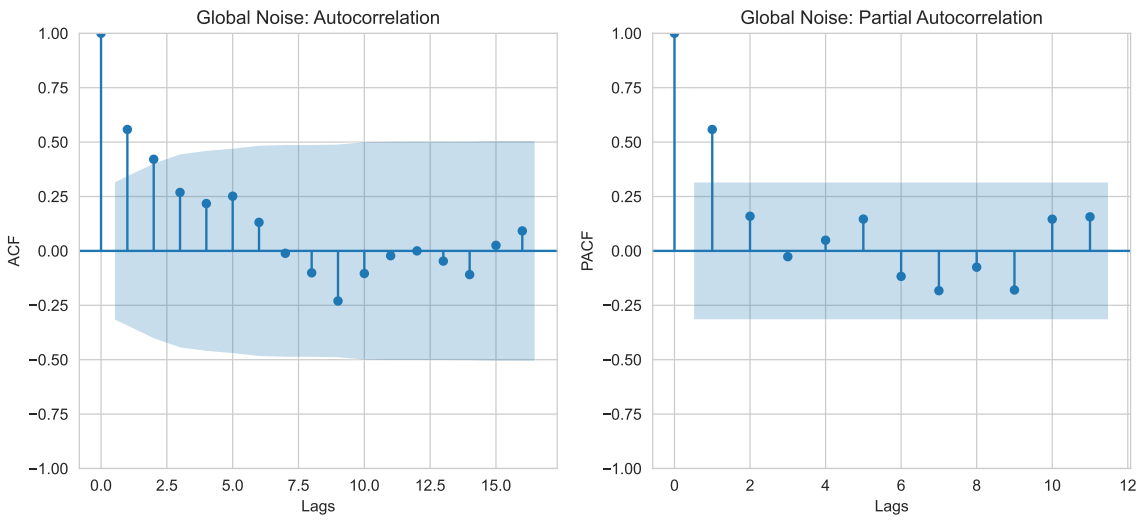

Figure S13: **ACF and PACF for the cross-sectional average wind BOS costs.** We consider the cross-sectional (national) average BOS costs for wind. **(left panel)** The ACF plot for the resulting time series. **(right panel)** The partial PACF plot for the resulting time series. The ACF is declining 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.

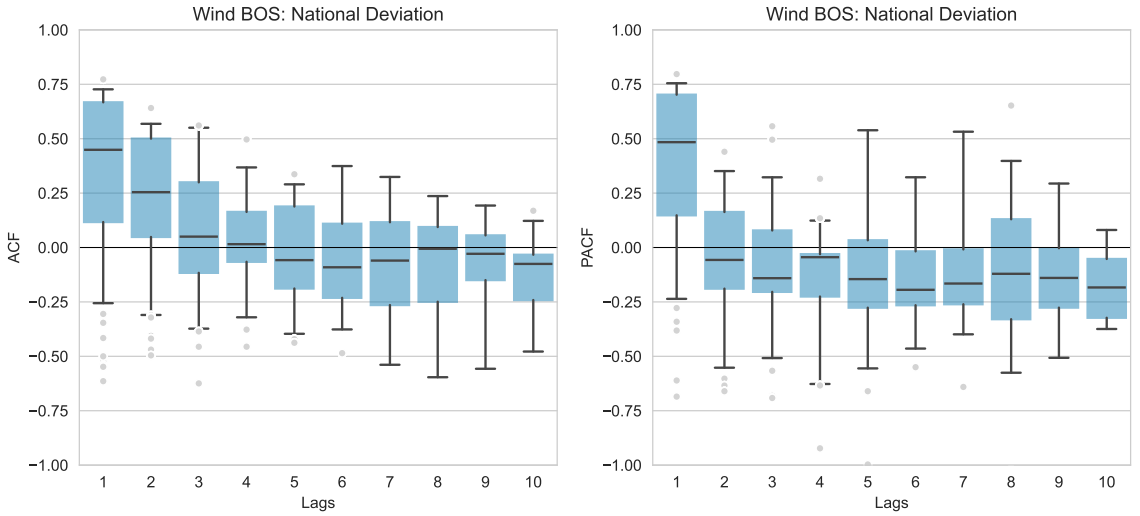

Figure S14: **Cross-sectional weighted average ACF and PACF for the national wind BOS cost deviations.** We consider the national BOS costs deviation from the global average. We calculate national ACF and PACF functions and take their cross-sectional averages. **(left panel)** The cross-sectional weighted 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.

Lag (Years)	0	0.5	1	1.5	2	2.5	3
Only Local learning ω	-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 S5: Regression results for wind BOS costs, according to Equation (S26), 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.

Table S5 shows our previous regression results for different lags. We focus on the 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.

S3.5.2 *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.[36] for fossil fuel prices and model future wind BOS costs as an AR(1) process around the historical average.

$$\log \mathcal{J}(t) = \phi \log \mathcal{J}(t-1) + \epsilon_t + \kappa, \tag{S28}$$

$$\log \mathcal{J}_j(t) - \log \mathcal{J}(t) \equiv \log \overline{\mathcal{J}}_j(t) = \varphi \log \overline{\mathcal{J}}_j(t-1) + \eta_{jt}. \tag{S29}$$

As for solar PV, Equation (S28) describes the global average wind BOS cost $\log \mathcal{J}(t)$. Equation (S29) describes the national deviation from the global average $\log \overline{\mathcal{J}}_j(t)$. Both the global average BOS costs (S28), as well as the national deviations thereof (S29) 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 (S28) and national deviations (S29) we need to determine the autocorrelation parameters ϕ and φ, 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 (S28) depends on the constant κ. For ϕ, we only have one (global) short time series available, making estimation very difficult. Therefore, we use the average value for fossil fuels in Way et al.[36] of $\phi = 0.88$. For φ, time series are also short, but we have multiple independent cross-sectional (national) time series. We can, therefore, estimate φ_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

$$\hat{\kappa} = \frac{1}{m} \sum_t \log \mathcal{J}(t)(1-\phi). \tag{S30}$$

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 ϵ and η based on the residuals of Equations (S28) and (S29).

As before, statistical significance is a challenge when it comes to validating this model due to the limited number of independent data points. For the global model (S28), an AR(1) process is the natural choice given its mean-reverting nature. Our choice of ϕ is also a natural choice that offers a good match to the empirically observed forecast accuracy. We focus our validation on the national deviation model (S29) where we can make statistically significant statements. We further do not include the estimation of φ 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 σ_η.[45]

Our procedure is the same as previously described; we consider 5-year moving average windows to estimate σ_η and make stochastic forecasts for the respective future of $\log \overline{\mathcal{J}}_j(t)$. We compare the distribution of our forecasts to that of the residuals across quantiles p using the cumulative PIT. We

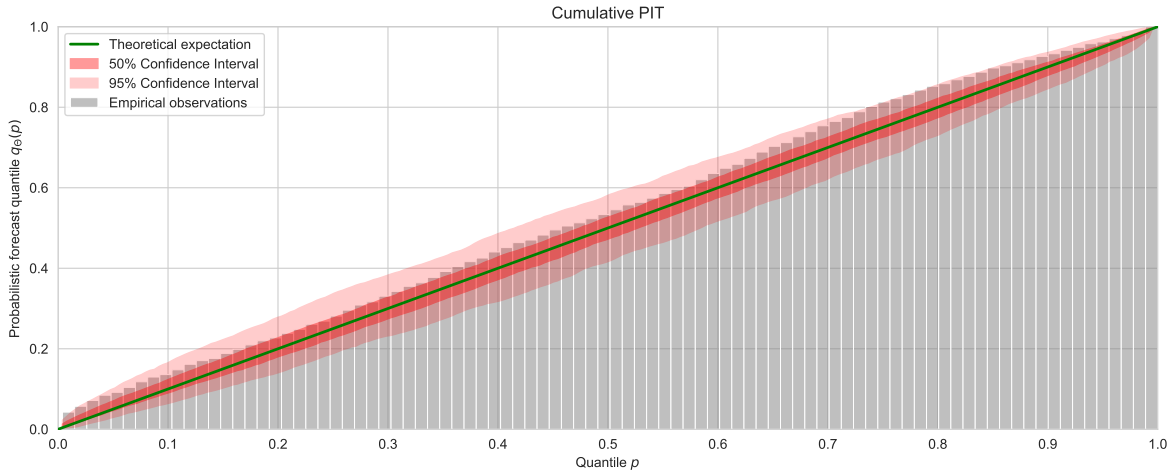

Figure S15: **Out-of-sample validation of the national wind BOS cost forecast error, relative to the true global average cost.** We consider only the national deviation of wind BOS cost relative to the global average BOS cost (S29). We estimate model parameters σ_η on 5-year moving average windows on 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 gray. The green 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.

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[4] 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 dependence of costs between these countries.

Figure S15 shows the observed cumulative PIT $q_\Theta(p)$ against p. The gray 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.

S3.6 Investment cost forecasts

Summary: *We aggregate our forecasts for the modules/turbines and BOS costs to obtain a forecast for the total investment cost. For wind, the difference in learning rates between turbines and BOS costs implies that our aggregate model for investment cost forecasts is preferred over a pure Wright's law model.For solar, 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.*

When we are interested in the total investment costs I, we can sum the disaggregate module/turbine and BOS forecasts according to

$$I_{jt} = \mathcal{M}_{jt} + \mathcal{J}_{jt}. \tag{S31}$$

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 I. However, if we have little data to estimate our model parameters, using Wright's law on the total investment cost I directly could be preferred over forecasting turbine/modules and BOS costs independently.[33,34,51,52] The reason for this is the bias-variance trade-off discussed earlier.[53] 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 I.

This topic warrants further study outside the scope of this work.[54] 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.

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 surrogate data for two time series X and Y. Both of them follow an equivalent process with different parameters, i.e.,

$$\text{solar module / wind turbine costs: } X_t = X_{t-1} \mathrm{e}^{\theta_X} \mathrm{e}^{\epsilon_t} \tag{S32}$$

$$\text{BOS costs: } Y_t = Y_{t-1} \mathrm{e}^{\theta_Y} \mathrm{e}^{\epsilon'_t}, \tag{S33}$$

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 $\theta_X = \theta_Y = -0.1$, close to what we observe for solar modules and BOS costs. In the second, we consider X evolving as a geometric random walk with $\theta_X = -0.1$ 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 make forecasts using both disaggregate and aggregate model parameters. We fit the forecast model on a subset of m years and 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 S16. 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.

S3.7 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,[4,55] while others describe capacity factors as historically constant.[56–58] We investigate trends in the capacity factor based on historical electricity generation and capacity.

S3.7.1 *Capacity Factor stationarity*

To assess the stationarity of a country's capacity factor, we follow the approach by Bolson et al.[59] 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 the Global Solar Atlas,[60] which

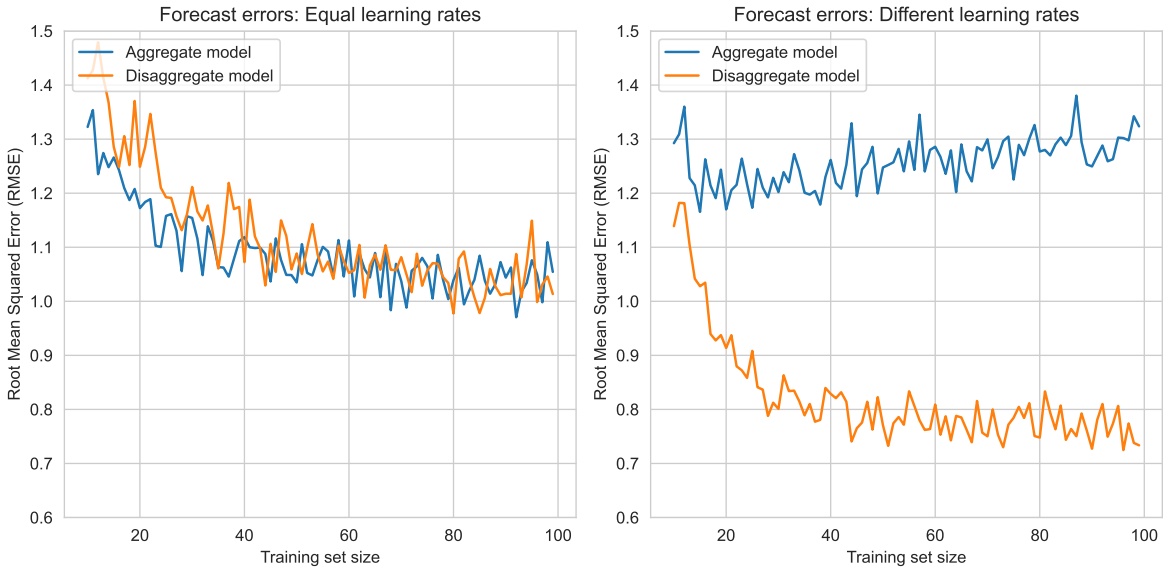

Figure S16: **Surrogate out-of-sample residuals for aggregate and disaggregate models, depending on the underlying processes.** We simulate the sum of two geometric random walks **(left, equivalent to solar)** and the sum of a geometric random walk and an AR(1) process **(right panel, equivalent to wind)**. 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.

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.[59]

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%.[61,62]

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

$$Z_j(t) = Z_j(t-1) + X_j(t), \tag{S34}$$

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

$$\Delta Y_j(t) \equiv Y_j(t) - Y_j(t-1) = \frac{1}{2}\left[\mathcal{Y}_j(t)X_j(t) + \mathcal{Y}_j(t-1)X_j(t-1)\right] - D_j(t). \tag{S35}$$

$\mathcal{Y}_j(t)X_j(t)$ 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

$$D_j(t) = \delta * Y_j(t), \tag{S36}$$

where we assume a fixed annual decline in power output δ. For wind, this degradation is around 1.6% annually,[61] while for solar, it is around 1.2% annually.[62] If we consider the weighted average capacity factor $\mathcal{Y}'_t$ for newly built resources in years t and $t-1$, as

$$\mathcal{Y}_j^{'}(t) = \frac{\mathcal{Y}_j(t)X_j(t) + \mathcal{Y}_j(t-1)X_j(t-1)}{X_j(t) + X_j(t-1)}. \tag{S37}$$

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

$$\mathcal{Y}_j^{'}(t) = \frac{2\Delta_1 Y_j(t)}{X_j(t) + X_j(t-1)} + \frac{2D_j(t)}{X_j^k(t) + X_j(t-1)}, \tag{S38}$$

where Δ_1 indicates that we are taking the difference between time t and $t-1$.

Based on this 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

$$\log \mathcal{Y}_j^{'}(t) = \log\left[\mathcal{Y}_j(t)X_j(t) + \mathcal{Y}_j(t-1)X_j(t-1)\right] - \log\left[X_j(t) + X_j(t-1)\right] + \epsilon_t + \rho_t * \epsilon_{t-1}, \tag{S39}$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma_t^2)$ and ρ_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 centered around 0, i.e.

$$H_0\colon\ \Delta_\tau \log \mathcal{Y}_j^{'}(t) \equiv \frac{1}{\tau}\left(\log \mathcal{Y}_j^{'}(t) - \log \mathcal{Y}_j^{'}(t+\tau)\right) \sim \mathcal{N}(0, \frac{\sigma_t^{'2} + \sigma_{t-2}^{'2}}{\tau^2}), \tag{S40}$$

where we discard the autocorrelation term ρ in place of σ'. 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 S17 shows the observed distributions. Visually, we find that both distributions are symmetric and centred around zero. Using the result of Bakinov & Székely,[63] 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 the capacity factor in some countries, there is no statistically significant global trend towards higher capacity factors. This means that for both solar and wind we should forecast future national capacity factors as trend-stationary. For solar, this further means that global changes in the inverter loading ratio (ILR)[64–66] have not had a significant impact on the capacity factor.

The global average capacity factor, as in the average across national capacity factors, may still appear to be increasing. This analysis shows that that is spurious. Any change in the global average capacity factor is most likely driven by sample bias. Over time, more national capacity factor data becomes available. When we add this to the global average, the sample average changes even though national capacity factors are stationary.

There are two important caveats to our statistical analysis. Firstly, some countries may still exhibit changes in capacity factor due to regulation or land ownership. If better geographical sites are made available to renewable sources, that could improve the capacity factor. Vice versa, as the best sites fill up, the capacity factor could decrease. Our analysis shows that neither of these effects have had a large impact on national capacity factors so far. Secondly, the growing use of tracking modules and

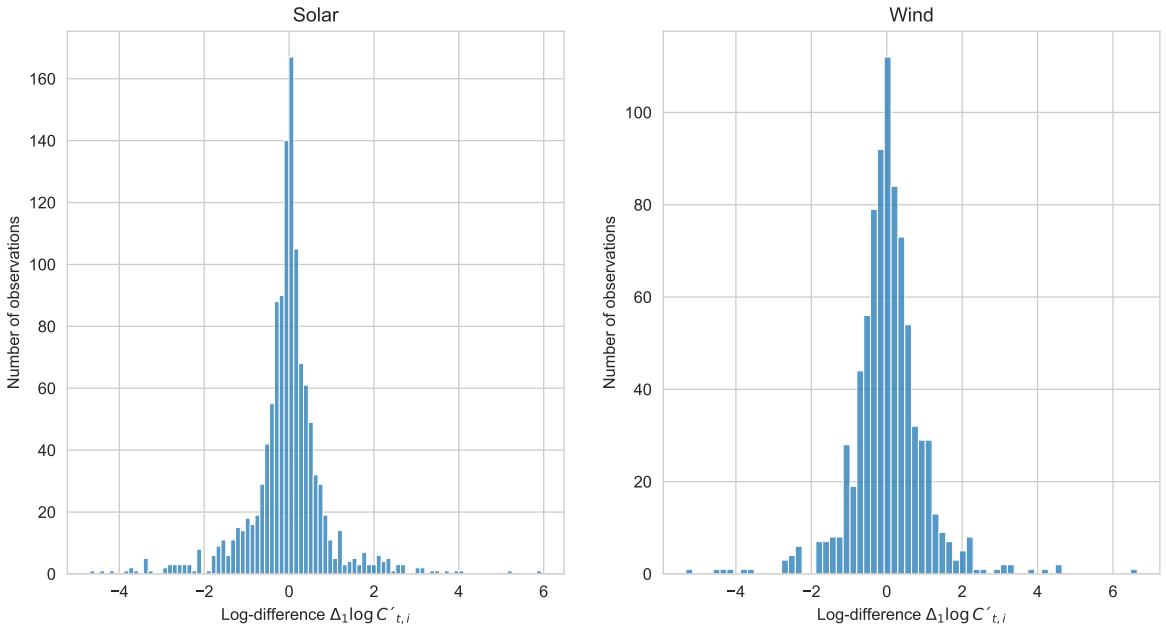

Figure S17: **Observed distribution for the normalized log-difference of the capacity factor for newly installed capacity.** Blue bars show the distribution of observed values for $\Delta_1 \log \mathcal{Y}'_j(t) = \log \mathcal{Y}'_j(t) - \log \mathcal{Y}'_j(t-1)$. For missing country-year data points, we take the next best higher-order difference in $\log \mathcal{Y}'_j(t)$ and normalize by the number of years included.

future use of bifacial and multi-junction modules may have some impact on the ILR and solar capacity factor.[67,68] These technologies allow solar assets to capture more sunlight over the course of a day, which could drive an increase in capacity factor. So far, this has not been the case but that might change if these technologies decrease in cost.

S3.7.2 *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

$$\text{Electricity Model: } \hat{\mathcal{Y}}_j = \frac{1}{m} \sum_{t=1}^{m} \frac{Y_j(t)}{Z_j(t)}, \tag{S41}$$

where m is the number of observations.[1]

For the LCOE-based model, we can use Equation (S12). We have seen in Section S3.1 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

[1] 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.

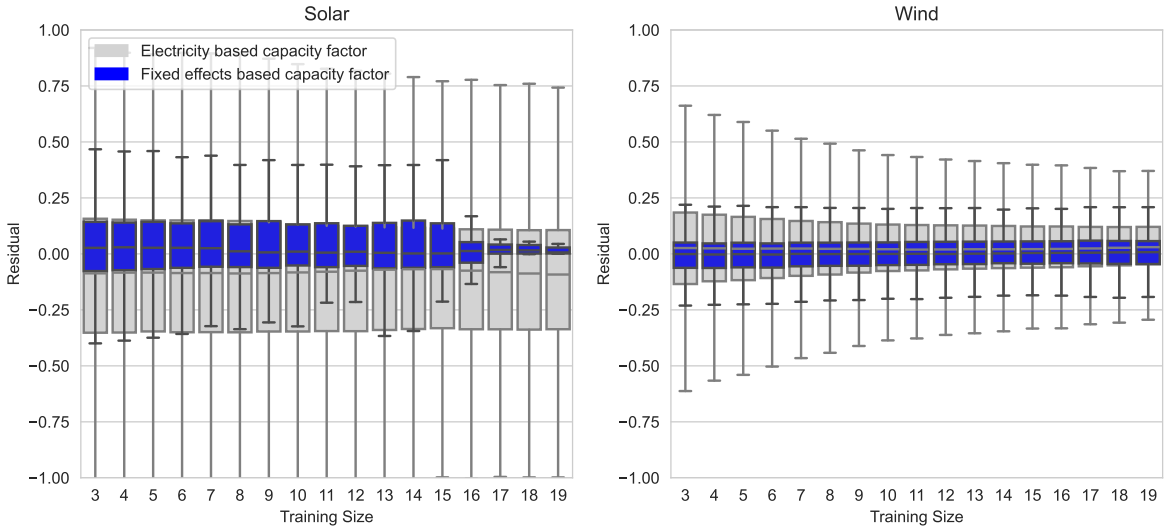

Figure S18: **Out-of-sample forecasts error for the capacity factor based on electricity generation and capacity, and a fixed-effect model for solar and wind.** We compare the out-of-sample forecast errors for the capacity factor in solar and wind. The x-axis shows the number of training years used in 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.

can also estimate the capacity factor as

$$\text{Fixed-effect Model:}\quad \log \hat{\mathcal{Y}}_j = \frac{1}{T_j}\sum_t \left(\log\left(I_j(t) * \mathcal{C}_j(t) + O_j(t)\right) - \log \mathcal{L}_j(t)\right) \tag{S42}$$

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:

$$\text{Observed data:}\quad \log \mathcal{Y}_j(t) = \log\left(I_j(t) * \mathcal{C}_j(t) + O_j(t)\right) - \log \mathcal{L}_j(t). \tag{S43}$$

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 S18. 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. 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.

S3.8 Operating & Maintenance cost forecasts

Summary: *Based on the limited data available, we show that operating & maintenance 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 S19 shows O&M costs against the investment costs. National O&M cost is very limited, with only two countries available for solar (Germany, US)[16,64] and six countries available for wind.[8,17,65] Statistical significance is therefore a key limitations.

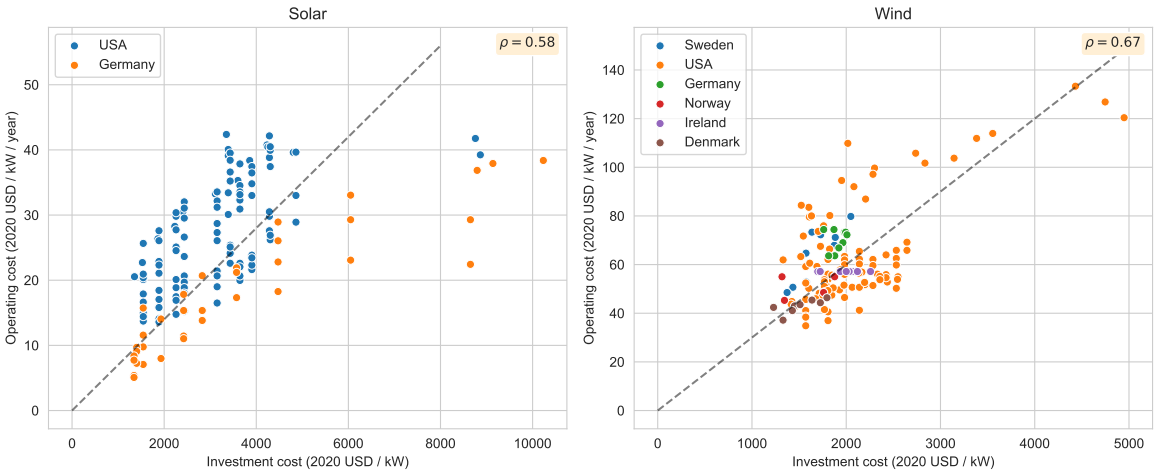

Figure S19: **Annual operating costs for wind and solar electricity vs investment costs.** Total investment costs are plotted against annual operating costs per kW capacity for solar **(left)** and wind **(right)**. The dotted lines indicate a fixed 3% relationship for wind and 1% for solar. Solar displays a Pearson correlation coefficient of ~ 0.58, wind ~ 0.67.[8,17,65]

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 ~ 0.58 for solar and ~ 0.67 for wind.

S3.8.1 *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 γ_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.

For wind, γ_j is very similar across countries. For solar, γ_j differs between Germany and the US. As evident in Figure S19, this is primarily driven by 8 US assets with investment costs exceeding USD 6 mn per kW. If we exclude this data, the γ_j values are again very similar. Given that solar investment costs are decreasing in the US and elsewhere, these data points are not representative of future O&M costs. Instead, we consider γ_j to be equivalent and stationary across countries and time for both solar and wind.

For our forecasts, we thus apply a fixed relative O&M cost $\gamma_i(t) = \gamma$ for wind and solar. This is consistent with the Steffen et al.[16] who estimate learning rates to forecast O& M costs. Here, O&M costs follow Wright's law, decreasing with cumulative electricity production at a similar rate to total investment costs. If both O&M costs O and total investment costs I decline at similar rates, the relative costs $\gamma = O/I$ is approximately constant.

S3.8.2 *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:

$$\log\frac{O_{ij}}{I_{ij}} \sim \mathcal{N}\left(\log\hat{\gamma}, \hat{\sigma}^2\left(\frac{1}{m}+1\right)\right), \tag{S44}$$

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 $\log \frac{O_{ij}}{I_{ij}}$:

$$\log \hat{\gamma} = \frac{1}{m} \sum_{i=1}^{m} \log \frac{O_{ij}}{I_{ij}}, \tag{S45}$$

$$\hat{\sigma}^2 = \frac{m^{k-1}}{m-1} \sum_{i=1}^{m} \left(\log \frac{O_{ij}}{I_{ij}} - \log \hat{\gamma} \right)^2. \tag{S46}$$

Our stochastic model of $\hat{\gamma}$ does not include the asset-specific exogenous shocks ϵ_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 estimate $\gamma \approx 3.1\%$ for wind and $\gamma \approx 0.8\%$, similar to earlier studies. We estimate the sample variance $\sigma \approx 0.13$ for solar and $\sigma \approx 0.05$ for wind.

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:

$$\mathcal{E}_i = \log \frac{O_{ij}}{I_{ij}} - \log \hat{\gamma}. \tag{S47}$$

Since our training-samples will have higher or lower variance $\hat{\sigma}$, we normalize the error by the estimated standard deviation

$$\frac{1}{\sqrt{\frac{1}{m} + 1}} \frac{\mathcal{E}_i}{\hat{\sigma}} \sim t(m-1). \tag{S48}$$

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 S20. We see that the t-distribution describes the out-of-sample error well and, thus, serves as an appropriate stochastic error model.

S4. 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.[69] In Section S6, we show additional country examples. In Section S4.3, we show that if scenarios are corrected for the initial value, the bias in their forecasts is reduced substantially.

S4.1 Scenario data collection

We construct an ensemble of scenarios from the AR6 database.[69] These scenarios are the ones that were considered for the sixth assessment report of the IPCC in 2022, although many scenarios were constructed in earlier years. 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).

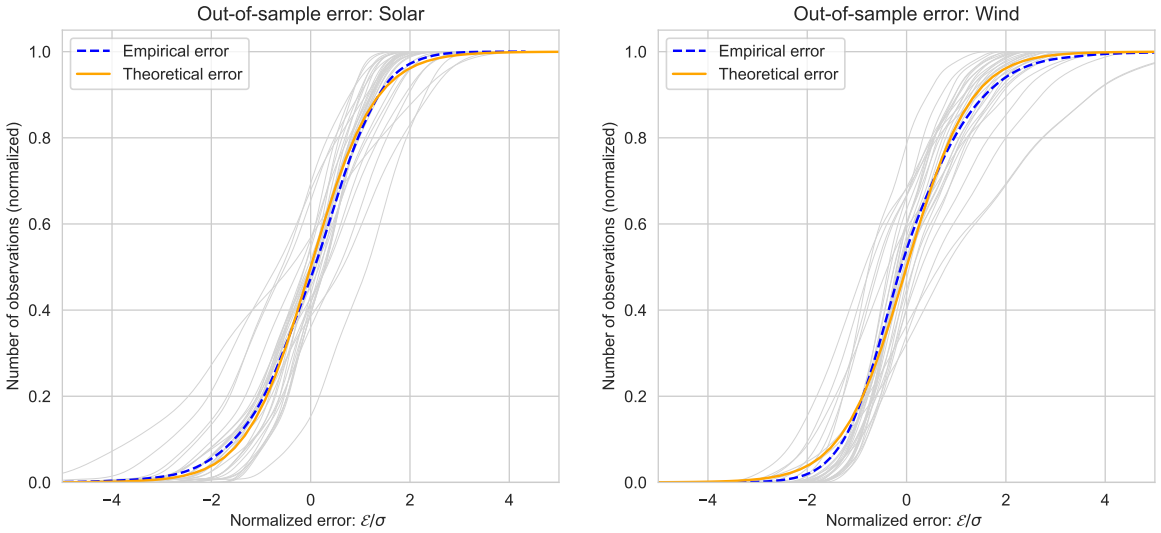

Figure S20: **Observed out of sample error of the relative O&M share $\log\gamma$ compared to the theoretical error.** The gray lines show the normalized out-of-sample errors of the observed relative O&M costs $\log\gamma$. The blue line shows the average over the gray lines. The orange line shows the theoretical distribution of the out-of-sample error, as per Equation (S3.8.2). 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.

S4.2 Comparison of scenarios and our total investment forecasts

As outlined in the main text, we compare the total investment cost I for solar and wind in 2040 in the AR6 scenarios to our forecasted values for I. For the scenario costs, we take the total investment costs in 2040 directly from the AR6 database. For our forecasts, we make forecasts for I conditionally on the global solar and wind capacity in 2040.

To make our forecasts, we perform three steps.

- Set the global capacity for solar and wind in 2040. We vary this number between the global capacity in 2020 (731 GW for wind, 713 GW for solar) and 25 TW for solar and 10 TW for wind.
- Interpolate linearly between the capacity in 2020 and 2040 to obtain a time series of future capacities.
- Conditional on the capacity build out between 2020 and 2040, make forecasts for the total investment cost using our model for turbines, modules and BOS costs.

This methods provides stochastic forecasts for the total investment costs in 2040, conditional on the global cumulative capacity.

As explained in the main text, we observe discrepancies between our forecasts and the IAM scenarios, where our median (point) forecasts is mostly below the IAM scenario costs.

S4.3 Initial value correction

The reason for this discrepancy are the initial values of the total investment cost and global capacity used in the IAM scenario construction. Many IAMs are quickly outdated due to the rapid growth of renewable capacity and associated drop in technology costs. As a result, their 2020 costs, as published in 2022, are almost all too high.

Figure S21 displays the reported solar and wind capacity for 2020 against the actual data. We see that in both cases, the investment costs assumed in the IAMS are generally higher than what actually occurred. For solar, Remind 2.1 is the only model that reports costs below what was observed.

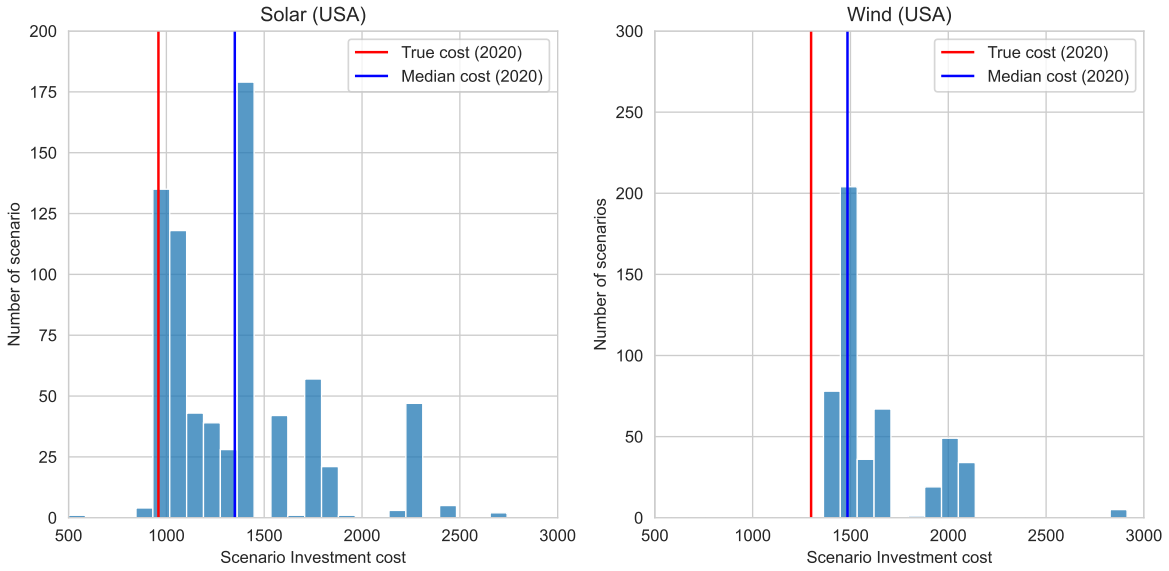

Figure S21: **Current solar and wind total investment cost for the US as reported in IAM scenarios** The blue bars show the solar (left) and wind (right) total investment costs in 2020, as reported by the AR6 scenarios with regionally differentiated capital costs. We exclude scenarios with costs exceeding 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.

We demonstrate that this initial value is the primary source of the difference between scenarios and our forecasts by re-scaling the IAM scenarios to the same initial values.

- For each scenario calculate the ratio of actual costs in 2020 and scenario costs in 2020,
- Multiply the scenario costs in 2040 with the resulting multiplicative factor,
- Repeat the previous comparison between our forecasts and the adjusted 2040 scenario costs.

The results of this for the US are shown in Figure S22. 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 2020 scenario data 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.,[70] and Way et al.[36]. Consequentially, our results emphasize the need for more stringent scenario validation and more frequent updates to the scenarios as the energy transition evolves.

S5. Estimated model parameters

The model parameters that were calibrated in this study are reported separately.

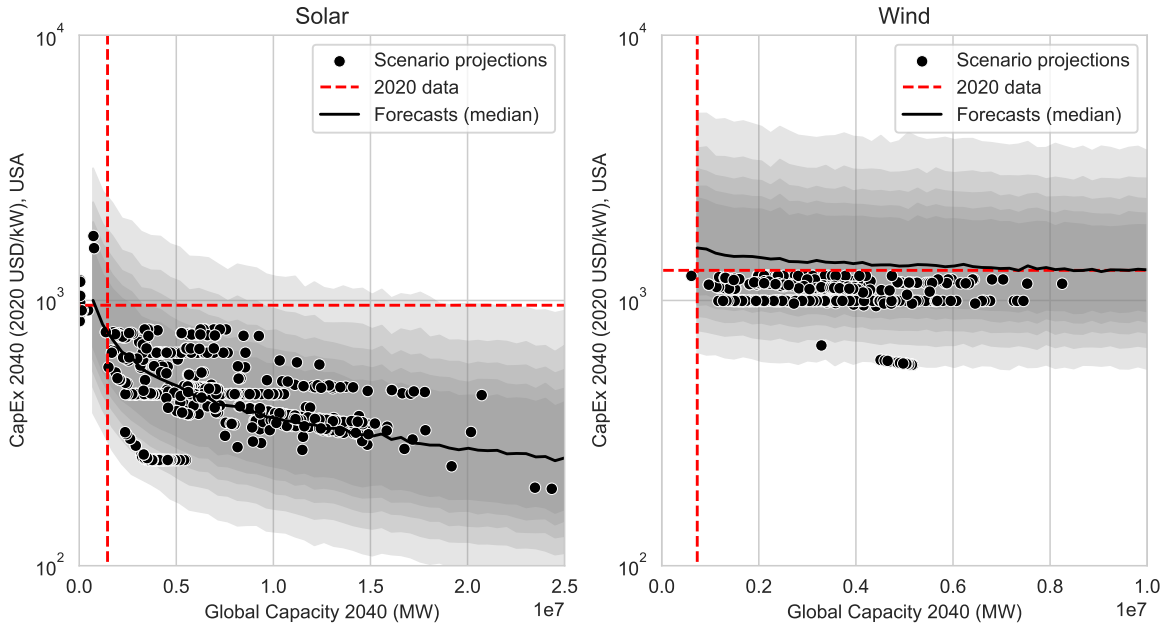

Figure S22: **Adjusted Solar and wind total investment cost projections reported in IAM scenarios against probabilistic projections from our model.** The dots show the total investment cost projections for the US in different scenarios, adjusted for different initial costs. The x-axis shows the associated 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.

S6. Additional Figures

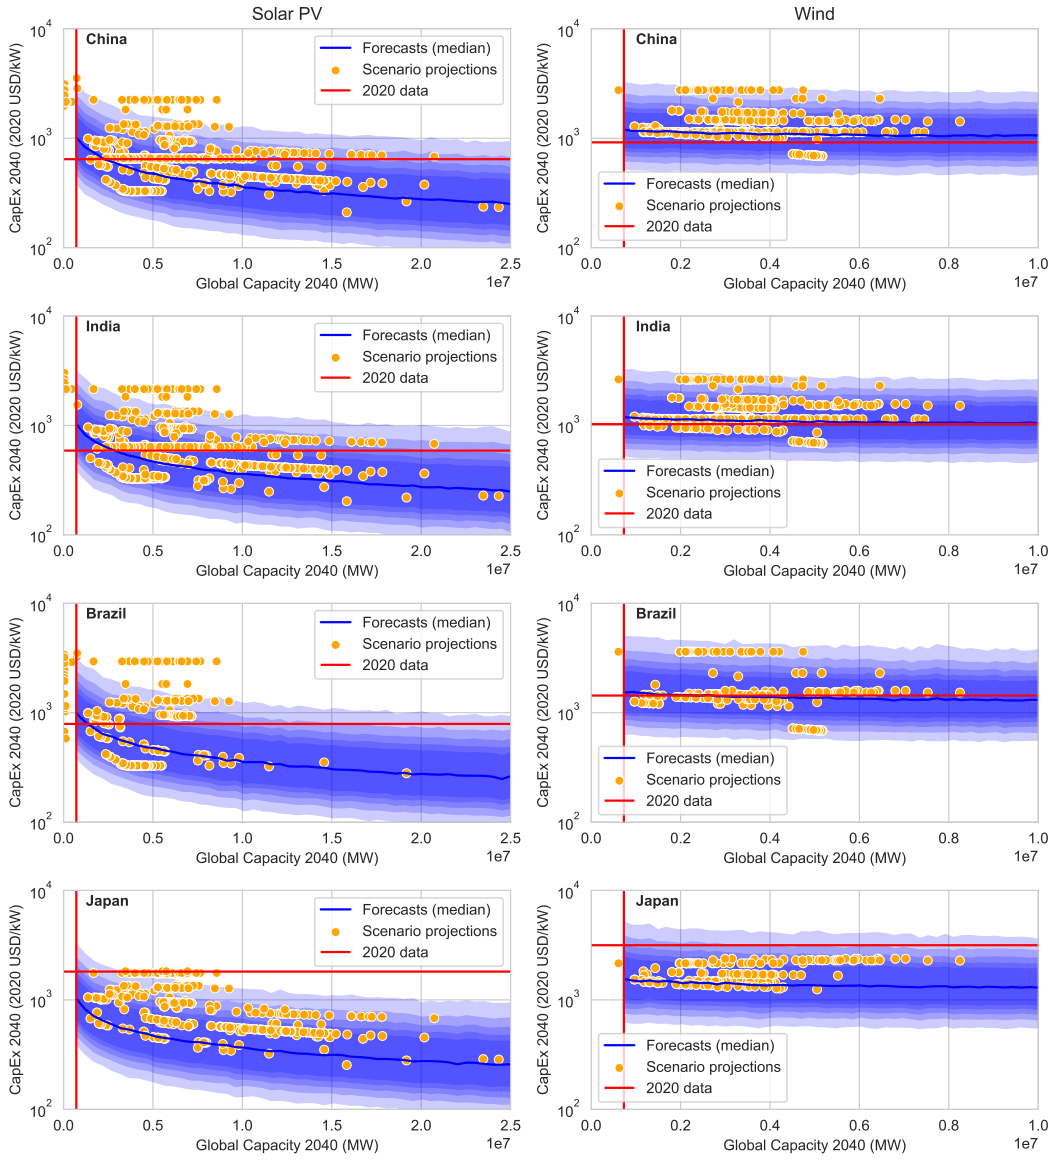

Figure S23: **Solar (left) and wind (right) total investment cost projections reported in IAM scenarios against probabilistic projections from our model. (not adjusted to initial cost and capacity)** The orange dots show the total investment cost projections for the different countries in different scenarios. 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.

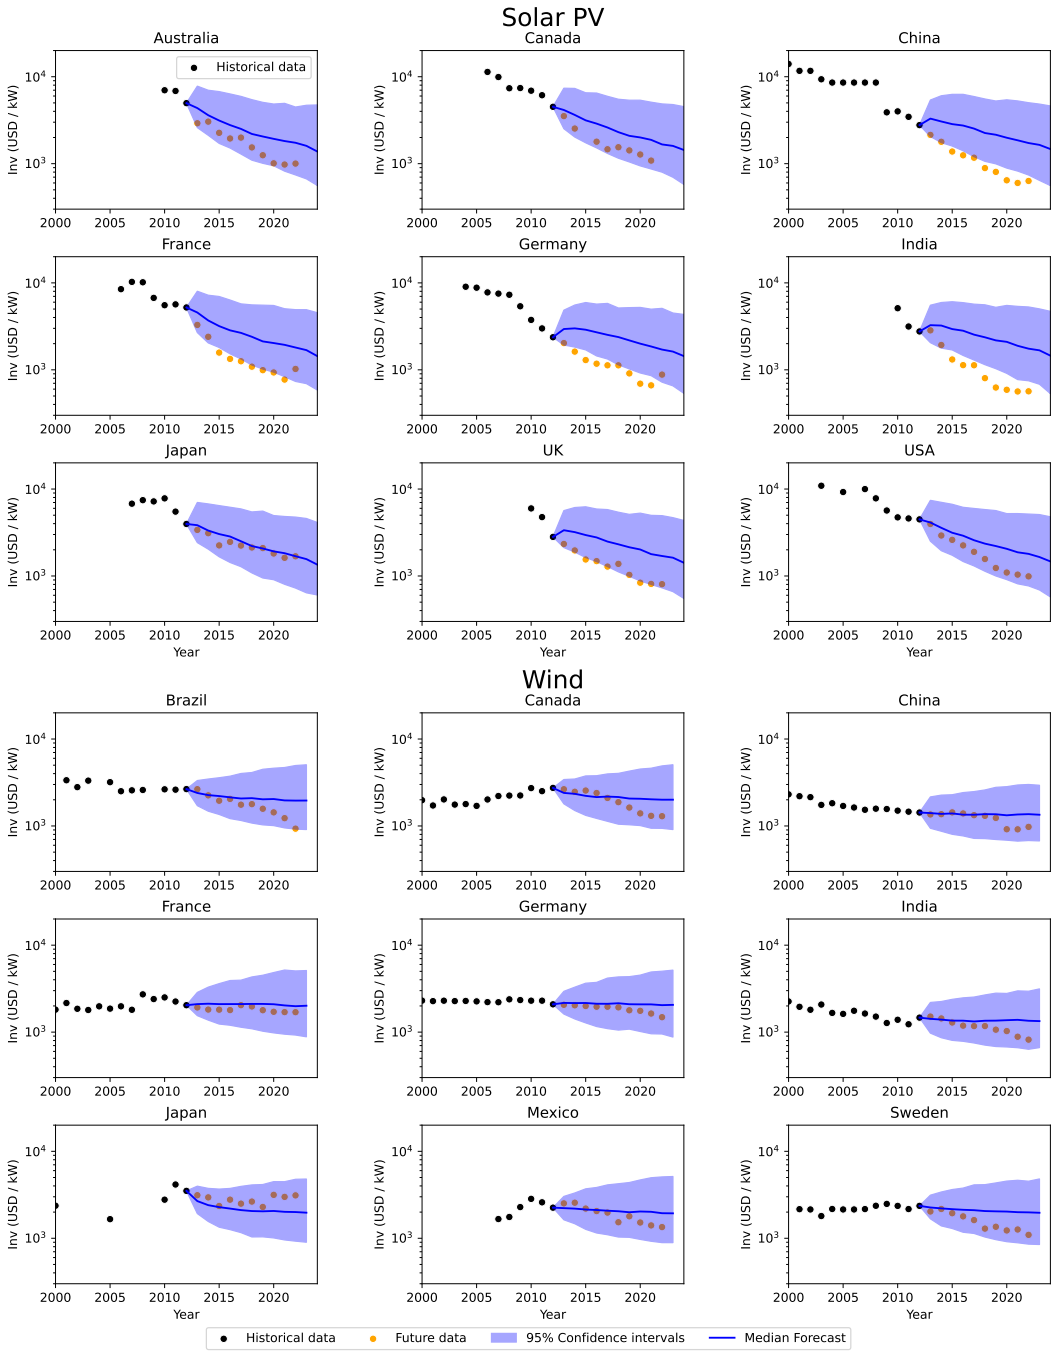

Figure S24: **Investment cost forecasts compared to true data.** We illustrate the Investment forecast model by comparing the stochastic model outcomes to the actual data for selected countries (in-sample). The top 9 panels show national forecasts for solar, and the bottom 9 panels show national forecasts 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.

Supplementary References

1. BP. *Statistical Review of World Energy* 2022. (2022).
2. IEA. *World Energy Balances* https://doi.org/10.1787/ data-00512-en. (2023).
3. IRENA. *Renewable Capacity Statistics 2022* tech. rep. (International Renewable Energy Agency, Abu Dhabi, Apr. 2022).
4. IRENA. *Renewable Power Generation Costs in 2021* tech. rep. (IRENA, 2021), 204.
5. IRENA. *The Cost of Financing for Renewable Power* tech. rep. (International Renewable Energy Agency, Abu Dhabi, 2023).
6. Elshurafa, A. M., Albardi, S. R., Bigerna, S. & Bollino, C. A. Estimating the Learning Curve of Solar PV Balance–of–System for over 20 Countries: Implications and Policy Recommendations. *Journal of Cleaner Production* **196,** 122–134. ISSN: 0959-6526. (2022) (Sept. 2018).
7. Kothari, S., Ameli, N. & Grubb, M. *Drivers of Balance-of-System Costs for Solar PV Technology* 2023.
8. Riva, A. D. *et al. 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* tech. rep. (National Renewable Energy Laboratory (NREL), Golden, CO (US), Nov. 2018).
9. IEA Photovoltaic Power Systems Program (PVPS). *Trends in Photovoltaic Applications 2022* tech. rep. 43:2022 (IEA PVPS TCP, 2022). (2023).
10. *Solar (Photovoltaic) Panel Prices* https://ourworldindata.org/grapher/solar-pv-prices. (2023).
11. Gosens, J. & Lu, Y. Prospects for Global Market Expansion of China's Wind Turbine Manufacturing Industry. *Energy Policy* **67,** 301–318. ISSN: 0301-4215 (Apr. 2014).
12. 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. (2024).
13. Beattie, S. *How Long Does It Take To Develop, Design, And Deliver A Commercial Solar PV Installation?* Oct. 2021. (2024).
14. Cruce, J., O'Shaughnessy, E., Harmon, J., Geiger, J. & Cook, J. *Residential Solar Adoption Timelines and Impacts from the COVID-19 Pandemic* tech. rep. NREL/TP-6A20-83529, 1889272, MainId:84302 (Sept. 2022), NREL/TP-6A20–83529, 1889272, MainId:84302. (2024).
15. U.S. Department of Energy. *Land-Based Wind Energy Economic Development Guide* tech. rep. (Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office: WindExchange, 2022). (2024).
16. Steffen, B., Beuse, M., Tautorat, P. & Schmidt, T. S. Experience Curves for Operations and Maintenance Costs of Renewable Energy Technologies. *Joule* **4,** 359–375. ISSN: 25424351. (2022) (Feb. 2020).
17. Wiser, R., Bolinger, M. & Lantz, E. Assessing Wind Power Operating Costs in the United States: Results from a Survey of Wind Industry Experts. *Renewable Energy Focus* **30,** 46–57. ISSN: 1755-0084. (2023) (Sept. 2019).
18. Wiser, R., Bolinger, M. & Seel, J. *Benchmarking Utility-Scale PV Operational Expenses and Project Lifetimes: Results from a Survey of U.S. Solar Industry Professionals* tech. rep. None, 1631678, ark:/13030/qt2pd8608q (June 2020), None, 1631678, ark:/13030/qt2pd8608q. (2023).
19. Comission de regulation de L'energie. Coûts et rentabilité des énergies renouvelables en France métropolitaine : Éolien terrestre, biomasse, solaire photovoltaïque (Apr. 2014).

20. Raupach-Sumiya, J., Matsubara, H., Prahl, A., Aretz, A. & Salecki, S. Regional Economic Effects of Renewable Energies - Comparing Germany and Japan. *Energy, Sustainability and Society* **5,** 10. ISSN: 2192-0567. (2023) (Dec. 2015).

21. Egli, F., Steffen, B. & Schmidt, T. S. A Dynamic Analysis of Financing Conditions for Renewable Energy Technologies. *Nature Energy* **3,** 1084–1092. ISSN: 20587546 (Dec. 2018).

22. Polzin, F. *et al.* The Effect of Differentiating Costs of Capital by Country and Technology on the European Energy Transition. *Climatic Change* **167.** ISSN: 15731480 (July 2021).

23. Bloomberg. *Bloomberg Terminal* Bloomberg Finance L.P. 2022.

24. Reuters. *Refinitiv Eikon* Reuters. 2022.

25. (IMF), I. M. F. *IMF Data* https://data.imf.org. 2022.

26. Kempa, K., Moslener, U. & Schenker, O. The Cost of Debt of Renewable and Non-Renewable Energy Firms. *Nature Energy* **6,** 135–142. ISSN: 2058-7546. (2023) (Feb. 2021).

27. Steffen, B. Estimating the Cost of Capital for Renewable Energy Projects. *Energy Economics* **88.** ISSN: 01409883 (May 2020).

28. Thornton, A. G. Africa Renewable Energy Discount Rate Survey – 2018 (2018).

29. Ondraczek, J., Komendantova, N. & Patt, A. *WACC the Dog: The Effect of Financing Costs on the Levelized Cost of Solar PV Power* May 2013.

30. Steffen, B. The Importance of Project Finance for Renewable Energy Projects. *Energy Economics* **69,** 280–294. ISSN: 01409883 (Jan. 2018).

31. Lafond, F. *et al.* How Well Do Experience Curves Predict Technological Progress? A Method for Making Distributional Forecasts. *Technological Forecasting and Social Change* **128,** 104–117. ISSN: 0040-1625 (Mar. 2018).

32. Barbose, G., Darghouth, N. R., O'Shaughnessy, E. & Forrester. *Tracking the Sun 2024 Edition* tech. rep. (Lawrence Berkley National Laboratory, Aug. 2024). (2025).

33. Grunfeld, Y. & Griliches, Z. Is Aggregation Necessarily Bad? *The Review of Economics and Statistics* **42,** 1–13. ISSN: 0034-6535. JSTOR: 1926089 (1960).

34. Lütkepohl, H. Forecasting Contemporaneously Aggregated Vector ARMA Processes. *Journal of Business & Economic Statistics* **2,** 201–214. ISSN: 0735-0015. JSTOR: 1391703 (1984).

35. PVPS, I. *Trends in Photovoltaic Applications 2015* tech. rep. 27:2015 (IEA PVPS TCP, 2015). (2023).

36. Way, R., Ives, M. C., Mealy, P. & Farmer, J. D. Empirically Grounded Technology Forecasts and the Energy Transition. *Joule* **6,** 1–26 (Sept. 2022).

37. IHSM. *Electric Plant Unit Wind Turbine Installed* Mar. 2023. (2023).

38. Hayashi, D., Huenteler, J. & Lewis, J. I. Gone with the Wind: A Learning Curve Analysis of China's Wind Power Industry. *Energy Policy* **120,** 38–51. ISSN: 0301-4215. (2023) (Sept. 2018).

39. Smart, B. & Farmer, J. D. *Macro and Micro Prediction Models* Unpublished. Oct. 2024.

40. Gneiting, T. & Katzfuss, M. Probabilistic Forecasting. *Annual Review of Statistics and Its Application* **1,** 125–151. ISSN: 2326-8298, 2326-831X. (2025) (Jan. 2014).

41. Gneiting, T., Balabdaoui, F. & Raftery, A. E. Probabilistic Forecasts, Calibration and Sharpness. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **69,** 243–268. ISSN: 1369-7412, 1467-9868. (2025) (Apr. 2007).

42. Pesaran, M. H. General Diagnostic Tests for Cross-Sectional Dependence in Panels. *Empirical Economics* **60,** 13–50. ISSN: 1435-8921. (2024) (Jan. 2021).

43. De Hoyos, R. E. & Sarafidis, V. Testing for Cross-Sectional Dependence in Panel-Data Models. *The Stata Journal* **6,** 482–496. ISSN: 1536-867X. (2024) (Nov. 2006).

44. Bretherton, C. S., Widmann, M., Dymnikov, V. P., Wallace, J. M. & Bladé, I. The Effective Number of Spatial Degrees of Freedom of a Time-Varying Field. *Journal of Climate* **12,** 1990–2009. ISSN: 1520-0442 (July 1999).

45. Farmer, J. D. & Lafond, F. How Predictable Is Technological Progress? *Research Policy* **45,** 647–665. ISSN: 00487333. arXiv: 1502.05274 (Apr. 2016).

46. Gneiting, T. & Raftery, A. E. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association* **102,** 359–378. ISSN: 0162-1459. (2024) (Mar. 2007).

47. Pesaran, M. H., Pick, A. & Timmermann, A. *Forecasting with Panel Data: Estimation Uncertainty versus Parameter Heterogeneity* Apr. 2024. arXiv: 2404.11198 [econ]. (2024).

48. Zamo, M. & Naveau, P. Estimation of the Continuous Ranked Probability Score with Limited Information and Applications to Ensemble Weather Forecasts. *Mathematical Geosciences* **50,** 209–234. ISSN: 1874-8953. (2024) (Feb. 2018).

49. Key, A., Robers, O. & Eberle, A. Scaling Trends for Balance-of-System Costs at Land-Based Wind Power Plants: Opportunities for Innovations in Foundation and Erection. *Wind Engineering* **46,** 896–913. (2023) (2022).

50. Pindyck, R. S. The Long-Run Evolution of Energy Prices. *The Energy Journal* **20,** 1–27. ISSN: 0195-6574. JSTOR: 41322828. (2024) (1999).

51. Bermingham, C. & D'Agostino, A. Understanding and Forecasting Aggregate and Disaggregate Price Dynamics. *Empirical Economics* **46,** 765–788. ISSN: 1435-8921. (2024) (Mar. 2014).

52. Hendry, D. F. & Hubrich, K. Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate. *Journal of Business & Economic Statistics* **29,** 216–227. ISSN: 0735-0015. JSTOR: 25800794 (2011).

53. Baumgärtner, C. L., Way, R., Ives, M. C. & Farmer, J. D. The Need for Better Statistical Testing in Data-Driven Energy Technology Modeling. *Joule* **8,** 2453–2466. ISSN: 2542-4351. (2024) (Sept. 2024).

54. Ferioli, F., Schoots, K. & van der Zwaan, B. Use and Limitations of Learning Curves for Energy Technology Policy: A Component-Learning Hypothesis. *Energy Policy* **37,** 2525–2535. ISSN: 03014215 (July 2009).

55. NREL. *Annual Technology Baseline 2022* https://atb.nrel.gov/. 2022. (2023).

56. Boretti, A. & Castelletto, S. Trends in Performance Factors of Wind Energy Facilities. *SN Applied Sciences* **2,** 1718. ISSN: 2523-3971. (2023) (Sept. 2020).

57. Bolinger, M. & Seel, J. *Utility-Scale Solar: Empirical Trends in Project Technology, Cost, Performance, and PPA Pricing in the United States – 2018 Edition* tech. rep. 1477381 (Sept. 2018), 1477381. (2023).

58. Boccard, N. Capacity Factor of Wind Power Realized Values vs. Estimates. *Energy Policy* **37,** 2679–2688. ISSN: 0301-4215. (2023) (July 2009).

59. Bolson, N., Prieto, P. & Patzek, T. Capacity Factors for Electrical Power Generation from Renewable and Nonrenewable Sources. *Proceedings of the National Academy of Sciences* **119,** e2205429119. (2023) (Dec. 2022).

60. ESMAP. *Global Solar Atlas* https://globalsolaratlas.info. (2023).

61. Staffell, I. & Green, R. How Does Wind Farm Performance Decline with Age? *Renewable Energy* **66,** 775–786. ISSN: 0960-1481. (2023) (June 2014).

62. Kim, J. *et al.* A Review of the Degradation of Photovoltaic Modules for Life Expectancy. *Energies* **14,** 4278. ISSN: 1996-1073. (2023) (Jan. 2021).

63. Bakirov, N. K. & Székely, G. J. Student's t-Test for Gaussian Scale Mixtures. *Journal of Mathematical Sciences* **139,** 6497–6505. ISSN: 1573-8795. (2023) (Dec. 2006).

64. IRENA. *Renewable Power Generation Costs in 2023* tech. rep. (IRENA, Abu Dhabi, 2024).

65. Bolinger, M., Seel, J., Warner, C. & Robson, D. *Utility-Scale Solar, 2022 Edition: Empirical Trends in Deployment, Technology, Cost, Performance, PPA Pricing, and Value in the United States* tech. rep. None, 1888246, ark:/13030/qt7496x1pc (Sept. 2022), None, 1888246, ark:/13030/qt7496x1pc. (2023).

66. Deschamps, E. M. & Rüther, R. Optimization of Inverter Loading Ratio for Grid Connected Photovoltaic Systems. *Solar Energy* **179,** 106–118. ISSN: 0038-092X. (2023) (Feb. 2019).

67. SolarWorld. *Calculating the Additional Energy Yield of Bifacial Solar Modules* tech. rep. (SolarWorld, Sept. 2021).

68. de Souza Silva, J. L. *et al. Impact of Bifacial Modules on the Inverter Clipping in Distributed Generation Photovoltaic Systems in Brazil* in *2021 Brazilian Power Electronics Conference (COBEP)* (Nov. 2021), 1–6.

69. Byers, E. *et al. AR6 Scenarios Database* https://zenodo.org/record/5886911. Nov. 2022.

70. Creutzig, F. *et al.* The Underestimated Potential of Solar Energy to Mitigate Climate Change. *Nature Energy* **2.** ISSN: 20587546. arXiv: 1710.05957 (Aug. 2017).