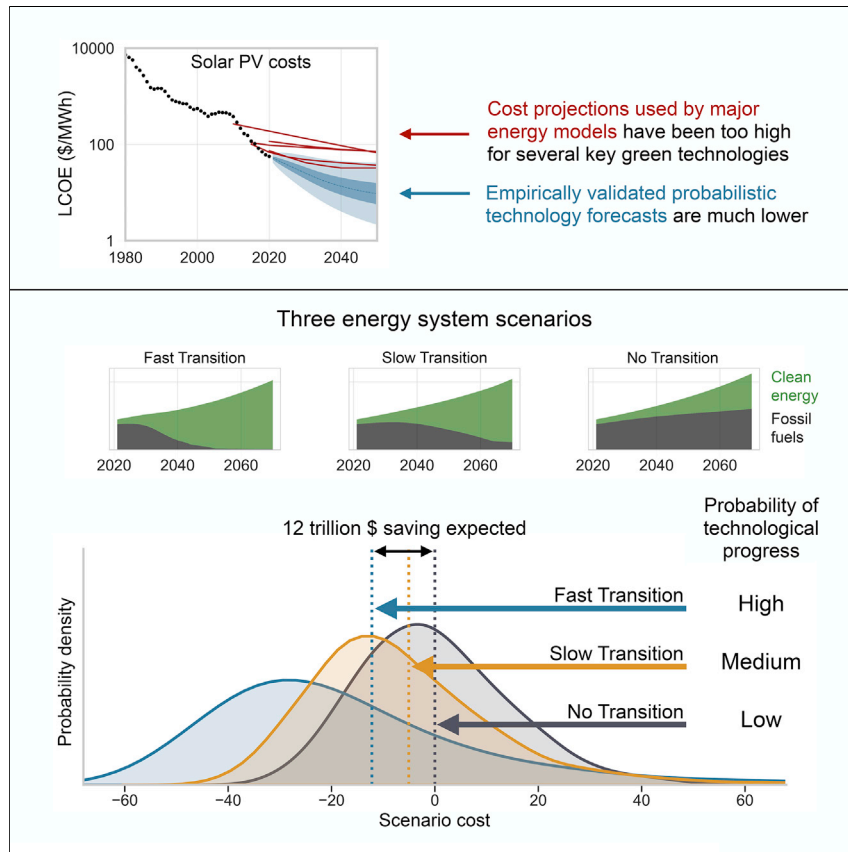


Article

Empirically grounded technology forecasts and the energy transition



Decisions about how and when to decarbonize the global energy system are highly influenced by estimates of the likely cost. Here, we generate empirically validated probabilistic forecasts of energy technology costs and use these to estimate future energy system costs under three scenarios. Compared to continuing with a fossil fuel-based system, a rapid green energy transition is likely to result in trillions of net savings, even without accounting for climate damages or climate policy co-benefits.

Rupert Way, Matthew C. Ives, Penny Mealy, J. Doayne Farmer

rupert.way@smithschool.ox.ac.uk

Highlights

Empirically validated probabilistic forecasts of energy technology costs

Future energy system costs are estimated for three different scenarios

A rapid green energy transition will likely result in trillions of net savings

Energy models should be updated to reflect high probability of low-cost renewables



Article

Empirically grounded technology forecasts and the energy transition

Rupert Way,^{1,2,6,*} Matthew C. Ives,^{1,2} Penny Mealy,^{1,2,3} and J. Doyne Farmer^{1,4,5}

SUMMARY

Rapidly decarbonizing the global energy system is critical for addressing climate change, but concerns about costs have been a barrier to implementation. Most energy-economy models have historically underestimated deployment rates for renewable energy technologies and overestimated their costs. These issues have driven calls for alternative approaches and more reliable technology forecasting methods. Here, we use an approach based on probabilistic cost forecasting methods that have been statistically validated by backtesting on more than 50 technologies. We generate probabilistic cost forecasts for solar energy, wind energy, batteries, and electrolyzers, conditional on deployment. We use these methods to estimate future energy system costs and explore how technology cost uncertainty propagates through to system costs in three different scenarios. Compared to continuing with a fossil fuel-based system, a rapid green energy transition will likely result in overall net savings of many trillions of dollars—even without accounting for climate damages or co-benefits of climate policy.

INTRODUCTION

Future energy system costs will be determined by a combination of technologies that produce, store, and distribute energy. Their costs and deployment will change with time due to innovation, competition, public policy, concerns about climate change, and other factors. To provide some perspective on the likely future energy system, Figure 1 shows how the energy landscape has evolved over the last 140 years. Figure 1A shows the historical costs of the principal energy technologies, and Figure 1B gives their deployment; both of which are on a logarithmic scale. As we approach the present in Figure 1A, the diagram becomes more congested, making it clear that we are in a period of unprecedented energy diversity, with many technologies with global average costs around \$100/MWh competing for dominance.

The long-term trends provide a clue as to how this competition may be resolved: The prices of fossil fuels such as coal, oil, and gas are volatile, but after adjusting for inflation, prices now are very similar to what they were 140 years ago, and there is no obvious long-range trend. In contrast, for several decades the costs of solar photovoltaics (PV), wind, and batteries have dropped (roughly) exponentially at a rate near 10% per year. The cost of solar PV has decreased by more than three orders of magnitude since its first commercial use in 1958.¹

Figure 1B shows how the use of technologies in the global energy landscape has evolved since 1880, when coal passed traditional biomass. It documents the slow exponential rise in the production of oil and natural gas over a century and the rapid rise and plateauing of nuclear energy. But perhaps the most remarkable feature is

CONTEXT & SCALE

Decisions about how and when to decarbonize the global energy system are highly influenced by estimates of the likely cost. Most energy-economy models have produced energy transition scenarios that overestimate costs due to underestimating renewable energy cost improvements and deployment rates. This paper generates probabilistic cost forecasts of energy technologies using a method that has been statistically validated on data for more than 50 technologies. Using this approach to estimate future energy system costs under three scenarios, we find that compared to continuing with a fossil fuel-based system, a rapid green energy transition is likely to result in trillions of net savings. Hence, even without accounting for climate damages or climate policy co-benefits, transitioning to a net-zero energy system by 2050 is likely to be economically beneficial. Updating models and expectations about transition costs could dramatically accelerate the decarbonization of global energy systems.



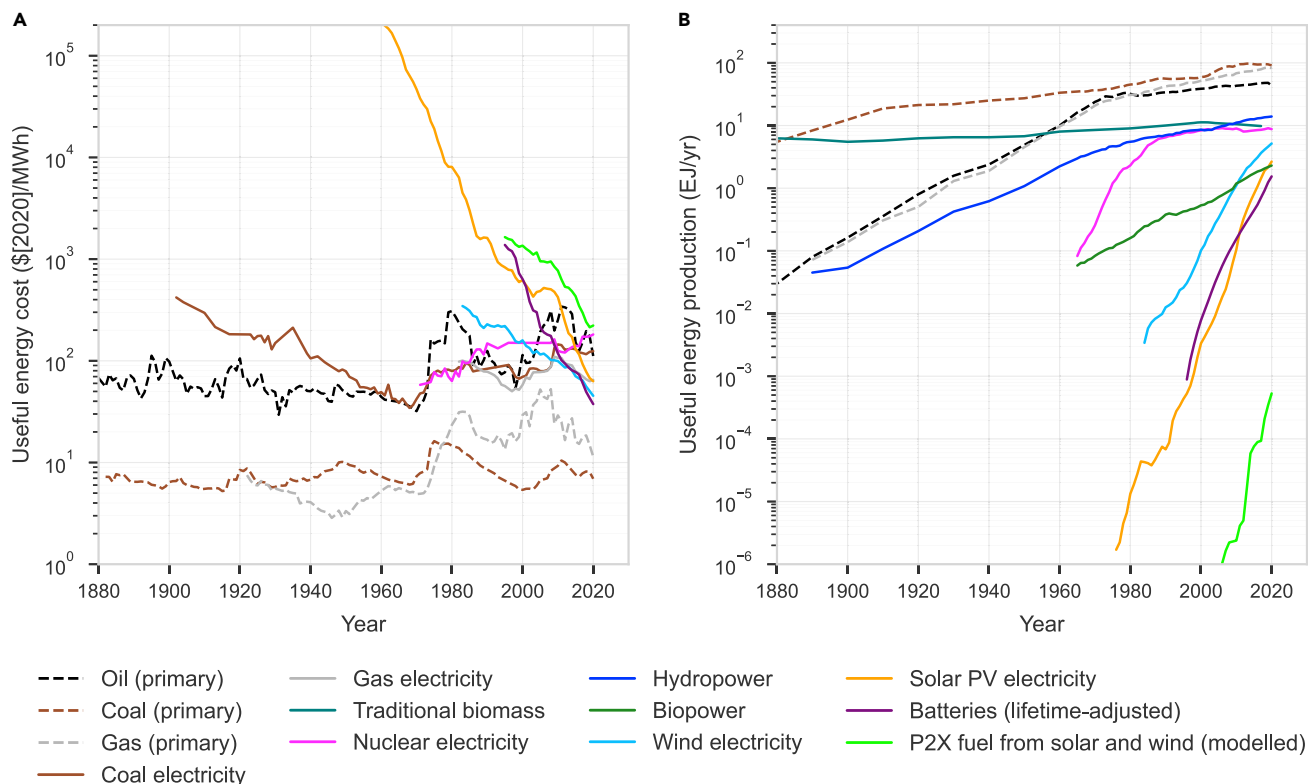


Figure 1. Historical costs and production of key energy supply technologies

(A) Inflation-adjusted useful energy costs (or prices for oil, coal, and gas) as a function of time. We show *useful energy* because it takes conversion efficiency into account (see Document S1 section “End-use conversion efficiencies”). Electricity generation technology costs are levelized costs of electricity (LCOEs). Battery series show capital cost per cycle and energy stored per year, assuming daily cycling for 10 years (these are not directly comparable with other data series here). Modeled costs of power-to-X (P2X) fuels, such as hydrogen or ammonia, assume historical polymer electrolyte membrane (PEM) electrolyzer costs and a 50–50 mix of solar and wind electricity.

(B) Global useful energy production. The provision of energy from solar photovoltaics has, on average, increased at 44% per year for the last 30 years, whereas wind has increased at 23% per year. These are just a few representative time series; all data sources and methods are given in Document S1 section “Data sources for Figure 1.”

the dramatic exponential rise in the deployment of solar PV, wind, batteries, and electrolyzers over the last decades as they transitioned from niche applications to mass markets. Their rate of increase is similar to that of nuclear energy in the 1970s, but unlike nuclear energy, they have all consistently experienced exponentially decreasing costs. The combination of exponentially decreasing costs and rapid exponentially increasing deployment is different from anything observed in any other energy technologies in the past, and positions these key green technologies to challenge the dominance of fossil fuels within a decade.

How likely is it that clean energy technology costs will continue to drop at similar rates in the future? Under what conditions will this happen, and what does this imply for the overall cost of the green energy transition? Is there a path forward that can get us to net-zero emissions cheaply and quickly? We address these questions here by applying empirically tested, state-of-the-art cost forecasting methods to energy technologies.

Historically, most energy-economy models have underestimated deployment rates for renewable energy technologies and overestimated their costs^{2–7}, which has led to calls for alternative approaches and more reliable technology forecasting

¹Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford OX1 3UQ, UK

²Smith School of Enterprise and the Environment, University of Oxford, Oxford OX1 3QY, UK

³SoDa Labs, Monash Business School, Monash University, Melbourne, VIC 3800, Australia

⁴Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

⁵Santa Fe Institute, Santa Fe, NM 87501, USA

⁶Lead contact

*Correspondence:
rupert.way@smithschool.ox.ac.uk
<https://doi.org/10.1016/j.joule.2022.08.009>

methods^{8–15}. Recent efforts have made progress in this direction^{16–19}, but they are largely deterministic in nature. The methods we use are probabilistic, allowing us to view energy pathways through the lens of placing bets on technologies. After all, powering modern economies requires betting on some technologies one way or another, be they clean technologies or more fossil fuels—the best we can do is make good bets. Which technologies should we bet on, and how likely are they to pay off? We focus on solar, wind, batteries, and electrolyzers, which we call here “key green technologies”, because they could play crucial roles in decarbonization and have strong progress trends that are well documented in publicly available datasets. We also consider the major incumbent energy technologies and compare our forecasts with projections made by influential energy-economy models. We investigate three different energy transition scenarios and discuss the implications for whole system costs and transition pathways.

Figure 1 provides a glimpse into the diverse nature of technological change as technologies rise and fall from dominance.^{20–22} It reflects how innovation and technological learning produce different outcomes for different technologies. The diversity of rates of technological improvement for energy technologies seen in Figure 1 is typical of technologies in general.^{23–25} Roughly speaking, technologies can be divided into two groups based on their rates of improvement. For the first group, comprising the vast majority of technologies, inflation-adjusted costs have remained roughly constant through time. Fossil fuels provide a good example: although there has been enormous progress in technologies for discovery and extraction, as easily accessible resources are depleted, it becomes necessary to extract less accessible resources, creating a “running-to-stand-still” dynamic in which prices have remained roughly constant for more than a century (this is true for all minerals^{26,27}). Another example of a non-improving technology is carbon capture and storage (CCS); despite significant effort, over its 50-year commercial history for enhanced oil recovery, costs have not declined at all.^{28,29} There are even cases, such as nuclear power, where costs have increased. By contrast, for a select group of improving technologies, costs have dropped roughly exponentially while deployment has increased exponentially.^{23–25} Rates of improvement for technologies such as optical fibers and transistors are as high as 40%–50% per year. Solar PV, wind, and batteries have behaved similarly but with improvement rates closer to 10% (see Document S1 section “The heterogeneity and persistence of technological change”). This makes unit costs for these technologies predictable, even if the specific technological innovations that lead to lower costs are not predictable.

Because the behavior of these two groups of technologies is so different, they require different cost forecasting models. Fossil fuels such as oil and gas are tradable commodities, and according to efficient markets theory, their prices should follow a random walk.³⁰ This provides a useful approximation for roughly a decade, but over longer spans of time they display mean reversion.^{31,32} This makes autoregressive models a natural choice, and we use them to forecast oil, coal, and gas prices (see [Experimental procedures](#) and Document S1 sections “AR(1) process,” “Oil,” “Coal,” and “Gas”).

For the select group of technologies that are improving, improvement rates are remarkably consistent.³³ For these technologies there are two dominant methods for quantitatively forecasting costs based on historical data. The first is a generalized form of *Moore’s law*, which says that costs drop exponentially as a function of time (i.e., at a fixed percentage per year).^{23,24,34} The second is *Wright’s law*, which

predicts that costs drop as a power law of cumulative production.³⁵ This relationship is also called an *experience curve* or *learning curve*, and cumulative production is also called *experience*. (For a discussion of challenges and caveats concerning Wright's law, see Document S1 section "Wright's law caveats.") Multifactor models have been proposed using additional input variables, such as patenting activity and research and development (R&D) expenditures, but data are limited and they require additional parameters. This can lead to overfitting, resulting in poor out-of-sample forecasts²⁵ (see Document S1 section "Bias-variance trade-off"). Multifactor models have so far not been properly tested, and we do not use them here.

Successful technologies tend to follow an "S-curve" for deployment, starting with a long phase of exponential growth in production that eventually tapers off due to market saturation.²² Under Wright's law, during the exponential growth phase costs drop exponentially in time, as they do for Moore's law, but when production growth eventually slows, cost improvement also slows. Improving technologies often spend many decades in the exponential growth phase, making it hard to distinguish between Moore's law and Wright's law. Forecasts using the two models have similar accuracy in backtesting experiments.²⁵

This brings up the important question of responsiveness to investment. Under Moore's law, costs are assumed to change exogenously over time, independent of policy and investment. Under Wright's law, costs depend on experience. Although experience does not directly cause costs to drop, it is correlated with other factors that do, such as level of effort and R&D, and has the essential advantage of being relatively easy to measure.^{36,37} For comparison, the historical time series displayed in [Figure 1](#) are plotted as experience curves in [Figure S17](#). The same heterogeneity of improvement rates seen in [Figure 1](#) is evident for Wright's law—the fact that fossil fuel prices have not dropped historically means that experience had no net effect—in stark contrast to key green technologies. In this paper, we focus on Wright's law because it satisfies the basic intuition that exerting greater effort induces greater effects. (We repeated all our modeling experiments using Moore's law and found that the qualitative conclusions are similar; see Document S1 section "Moore's law results." For a more thorough discussion of causality, see Document S1 section "Discussion on questions of causality.")

Wright's law has usually been used to generate point forecasts, meaning that the forecast is a deterministic function of experience, with no estimate of the error of the forecast. Early attempts at introducing error bars did not provide a *a priori* functional forms, which made the data requirements for out-of-sample testing prohibitive.^{25,38} More recently, *a priori* error estimates were derived that predict forecasting accuracy as a function of historical improvement rates and volatility, and the number of data points available for forecasting.^{33,39} Based on comprehensive backtesting, this method was shown to generate reliable *probabilistic* estimates of future costs. This was done by selecting reference dates in the past and then, using only the data available at the time, making forecasts over all time horizons up to 20 years into the future with respect to each reference date. Using historical data for 50 different technologies, based on roughly 6,000 forecasts, the empirically observed forecast accuracy closely matched the *a priori* derived estimates on all time horizons up to 20 years ahead.^{33,39} Our main contribution in this paper is to systematically apply this method—which we call the stochastic experience curve or stochastic Wright's law—to the energy transition.

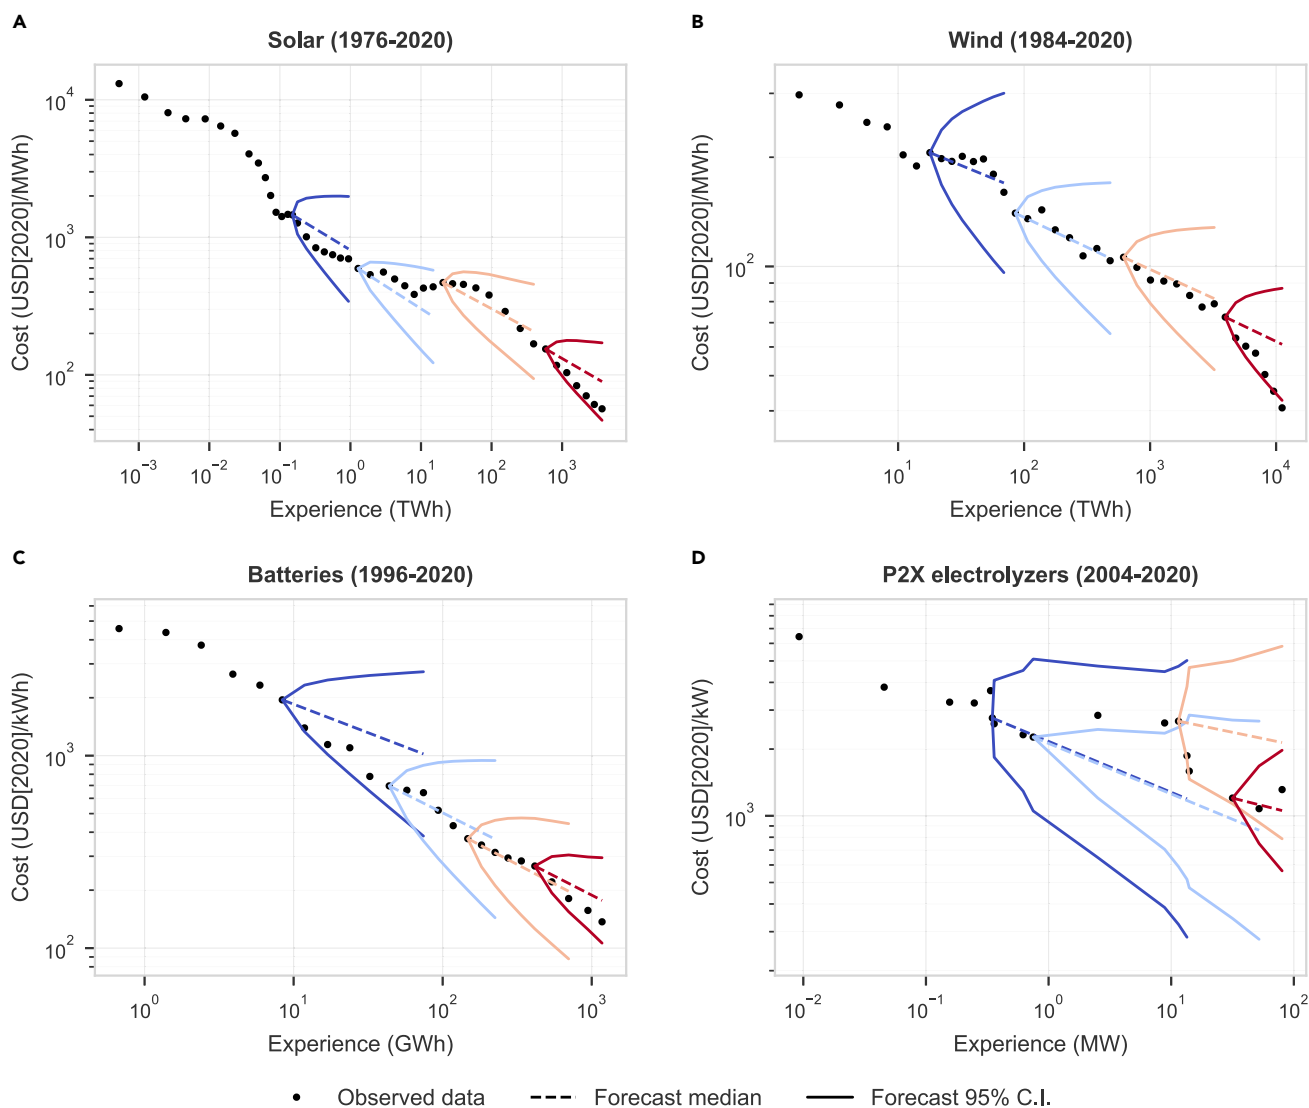


Figure 2. Historical performance of the stochastic experience curve forecasting method

(A–D) The four panels show stochastic Wright’s law applied to observed data for (A) solar, (B) wind, (C) batteries, and (D) P2X electrolyzers. Forecasts are made at regular intervals, using prior cost and deployment data to calibrate the model and “future” deployment data to generate the forecasts. Forecast medians and 95% confidence intervals (CIs) are shown, and colors denote forecast year, from earliest (dark blue) to most recent (red). Costs are LCOEs for solar and wind, and capacity costs for batteries and electrolyzers. P2X electrolyzers are assumed to be PEM electrolyzers here. See Document S1 section “Data, calibration and technology forecasts” for further details and data sources.

RESULTS

To test the accuracy of the stochastic experience curve method for forecasting costs of energy technologies, we applied it to historical data for solar, wind, batteries, and polymer electrolyte membrane (PEM) electrolyzers; the results are shown in Figure 2. Data prior to each forecast year were used to estimate parameters, then observed deployment data in subsequent years were used to generate forecasts conditioned on experience. The forecasts for solar, wind, and batteries are reasonable: most of the future values lie within the 95% confidence interval (CI), consistent with the *a priori* error estimates. As expected, forecast uncertainty decreases for later forecasts with more historical data. Due to the short dataset and high historical volatility, forecasts for electrolyzers are not as accurate, but the confidence intervals capture this.

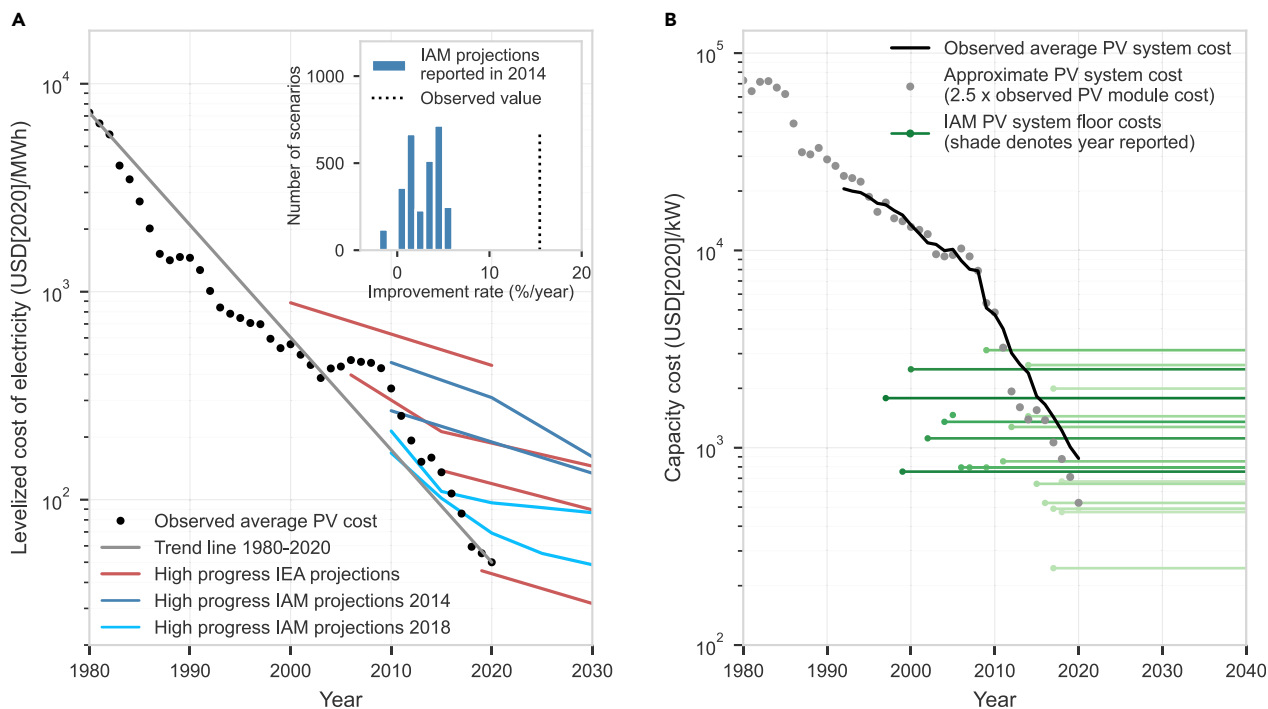


Figure 3. Historical PV cost projections and floor costs

(A) The black dots show the observed global average levelized cost of electricity (LCOE) over time. Red lines are LCOE projections reported by the International Energy Agency (IEA);⁸¹ dark blue lines are integrated assessment model (IAM) LCOE projections reported in 2014;⁴¹ and light blue lines are IAM projections reported in 2018.^{42,43} IAM projections are rooted in 2010 despite being produced in later years. The projections shown are exclusively “high technological progress” cost trajectories drawn from the most aggressive mitigation scenarios, corresponding to the largest projected cost reductions used in these models. Other projections made were even more pessimistic about future PV costs. The inset compares a histogram of projected compound annual reduction rates of PV system investment costs from 2010 to 2020 with what actually occurred (based on all 2,905 scenarios for which the data are available⁴¹).

(B) PV system floor costs implemented in a wide range of IAMs. The colors denote the year the floor cost was reported, ranging from 1997 (dark green) to 2020 (light green). Observed PV system costs are also shown. The cost of PV modules scaled by a constant factor of 2.5 is provided as a reference. For further details and data sources, see Figures 8 and 9A and Document S1 section “Solar PV electricity.”

It is instructive to compare the accuracy of the forecasts in Figure 2 to the outputs of influential global energy-economy models that are used to inform the Intergovernmental Panel on Climate Change (IPCC) and guide global climate policy.⁴⁰ Integrated assessment models (IAMs) are used to evaluate policies and generate scenarios for deployment and cost that are consistent with given climate targets under the assumption of optimal decision-making by economic agents. Their outputs are typically called “projections” to indicate that they are not intended to be used as forecasts. Figure 3 emphasizes this point. Figure 3A shows past projections of solar PV energy costs by the International Energy Agency’s (IEA) World Energy Model and several IAMs and compares them with the observed data.

The projections shown correspond to scenarios with the most aggressive climate policies and highest rates of technological innovation, i.e., those that produce the highest rates of key green technology deployment and the most optimistic cost declines. Nonetheless, their projected costs have been consistently much higher than historical trends. The inset of Figure 3A gives a histogram of all 2,905 projections of the annual rate at which solar PV system investment costs would fall between 2010 and 2020, as reported by nine separate IAM teams in the AMPERE modeling comparison project.⁴¹ The mean value of these projected cost reductions was 2.6%,

and all were less than 6%. In stark contrast, during this period, solar PV costs actually fell by 15% per year.

This makes it clear that it would have been a bad idea to treat these projections as conditional forecasts. By contrast, the stochastic experience curve method produces reliable conditional forecasts of known accuracy (and a published forecast of 2020 solar costs, made in 2010 using the deterministic version of Wright's law, was indeed far more accurate than any of the IAM or IEA projections made at the time⁴⁴). One of our goals in this paper is to illustrate how such forecasts are useful for planning the energy transition. (Note that IAM and IEA projections are better for mature incumbent technologies such as fossil fuels, but their projections for solar PV, wind, batteries, and electrolyzers have systematically underestimated deployment and overestimated costs.).^{2,5,45}

Wright's law is widely used to generate technology cost projections in IAMs.^{46–48} However, it is typically used in conjunction with ad hoc constraints such as deployment rate limits and floor costs, i.e., fixed levels that costs are assumed to never fall below. Because IAMs use costs to determine deployment (and vice versa), and many allow perfect foresight, constraints are necessary to prevent sharp cost declines due to Wright's law from leading to solutions in which key green technologies are deployed faster than is physically or socially plausible. It is difficult to know what constraints are realistic, which leads to ad hoc choices that strongly influence the results.

The historical record indicates that the constraints on key green technology deployment and costs used in IAMs have so far been much too stringent. For example, as shown in [Figure 3B](#), past floor costs used in IAMs have already repeatedly been violated. We know of no empirical evidence supporting floor costs and do not impose them (see Document S1 section "The use of floor costs in endogenous technological learning models"). Similarly, while there are likely limits to how quickly we can deploy key green technologies, it is difficult to know what they are. The outputs of IAMs depend critically on these constraints, which always alter the projections in the same direction, making them more pessimistic about the costs and deployment of key green technologies. As demonstrated here, the exponential growth of key green technologies and the relative accuracy of the unconstrained version of Wright's law suggest that thus far these constraints have not been binding. The imposition of excessively strong constraints is likely an important reason why the projections of these models have not corresponded to the historical record.

Probabilistic cost forecasts for individual technologies

We applied the methods discussed so far to make forecasts of future energy costs and prices. Given the reasons discussed in the introduction, for solar, wind, batteries, and electrolyzers, we used the stochastic form of Wright's law; and for oil, coal, and gas, we used an autoregressive model of order 1 (AR(1)). To generate experience curve forecasts, parameters for each technology were estimated from historical data. We then specified scenarios for the future deployment of each technology as a function of time and predicted a distribution of future costs.

We defined three representative deployment scenarios that we will explain in more detail later. The first scenario is consistent with the energy system transitioning away from fossil fuels by around 2050, and so we label this deployment scenario the "Fast Transition". The second scenario is consistent with eliminating fossil fuels by around 2070, so we label it the "Slow Transition". The final scenario is consistent with fossil

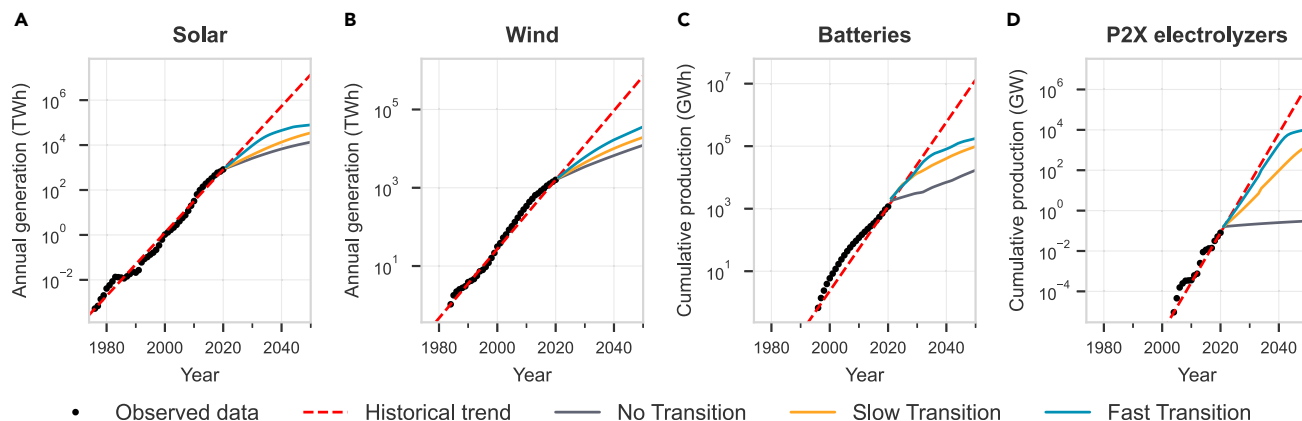


Figure 4. Technology deployment of key green technologies

Observed data up to 2020 are shown, plus three hypothetical growth scenarios up to 2050, corresponding to Fast Transition (blue), Slow Transition (yellow), and No Transition (gray). The trend line shown is the line through the first and last data points plotted and is indicative of the long-run trend so far.

fuels continuing to dominate the energy system, so we label it the “No Transition”. Figure 4 shows these three deployment scenarios for each of the four key green technologies in the context of their historical long-term growth trends. The deployment trajectories for these scenarios are consistent with the S-shaped curve widely observed in technology diffusion—the differences between them reflect differences in the timing and abruptness with which the growth of each technology tapers off.²²

Figure 5 shows probabilistic forecasts for seven particularly important energy technologies. The main panels of Figures 5A–5D show forecasts for key green technologies in the Fast Transition scenario, which are made using the stochastic version of Wright’s law. The insets show costs versus experience and emphasize that median costs develop identically as a function of experience in all scenarios. The side panels of Figures 5A–5D illustrate that under Wright’s law, forecast distributions depend on the scenario; as a result, in a faster transition, we are likely to reach lower costs sooner. Each Wright’s law technology initially follows its current trend of exponentially decreasing costs, but then progress slows as its rate of deployment drops. To generate fossil fuel cost forecasts, the AR(1) model was calibrated to observed data. For fossil fuels, model parameters depend on past data, but forecasts are independent of deployment, so each technology has a single forecast in all scenarios.

Figure 5 also shows a selection of future cost projections reported by IAM and IEA studies. As before, we show only their most optimistic projections, i.e., low cost projections that correspond to high technological progress scenarios. Consistent with the historical behavior of these models illustrated in Figure 3, the cost projections are high relative to historical trends. They are also all substantially higher than our forecast medians.

Of course, the deployment corresponding to these cost projections is not the same as that used to make our forecasts, so they are not perfectly comparable. However, as the boxplot panels show, the disparities persist across all our scenarios, including the No Transition scenario. This makes it clear that our cost forecasts are, all things equal, significantly lower than those used in these highly influential energy-economy models.

The stochastic version of Wright’s law we use here captures the historical volatility of technological progress and the associated parameter estimation error, and it

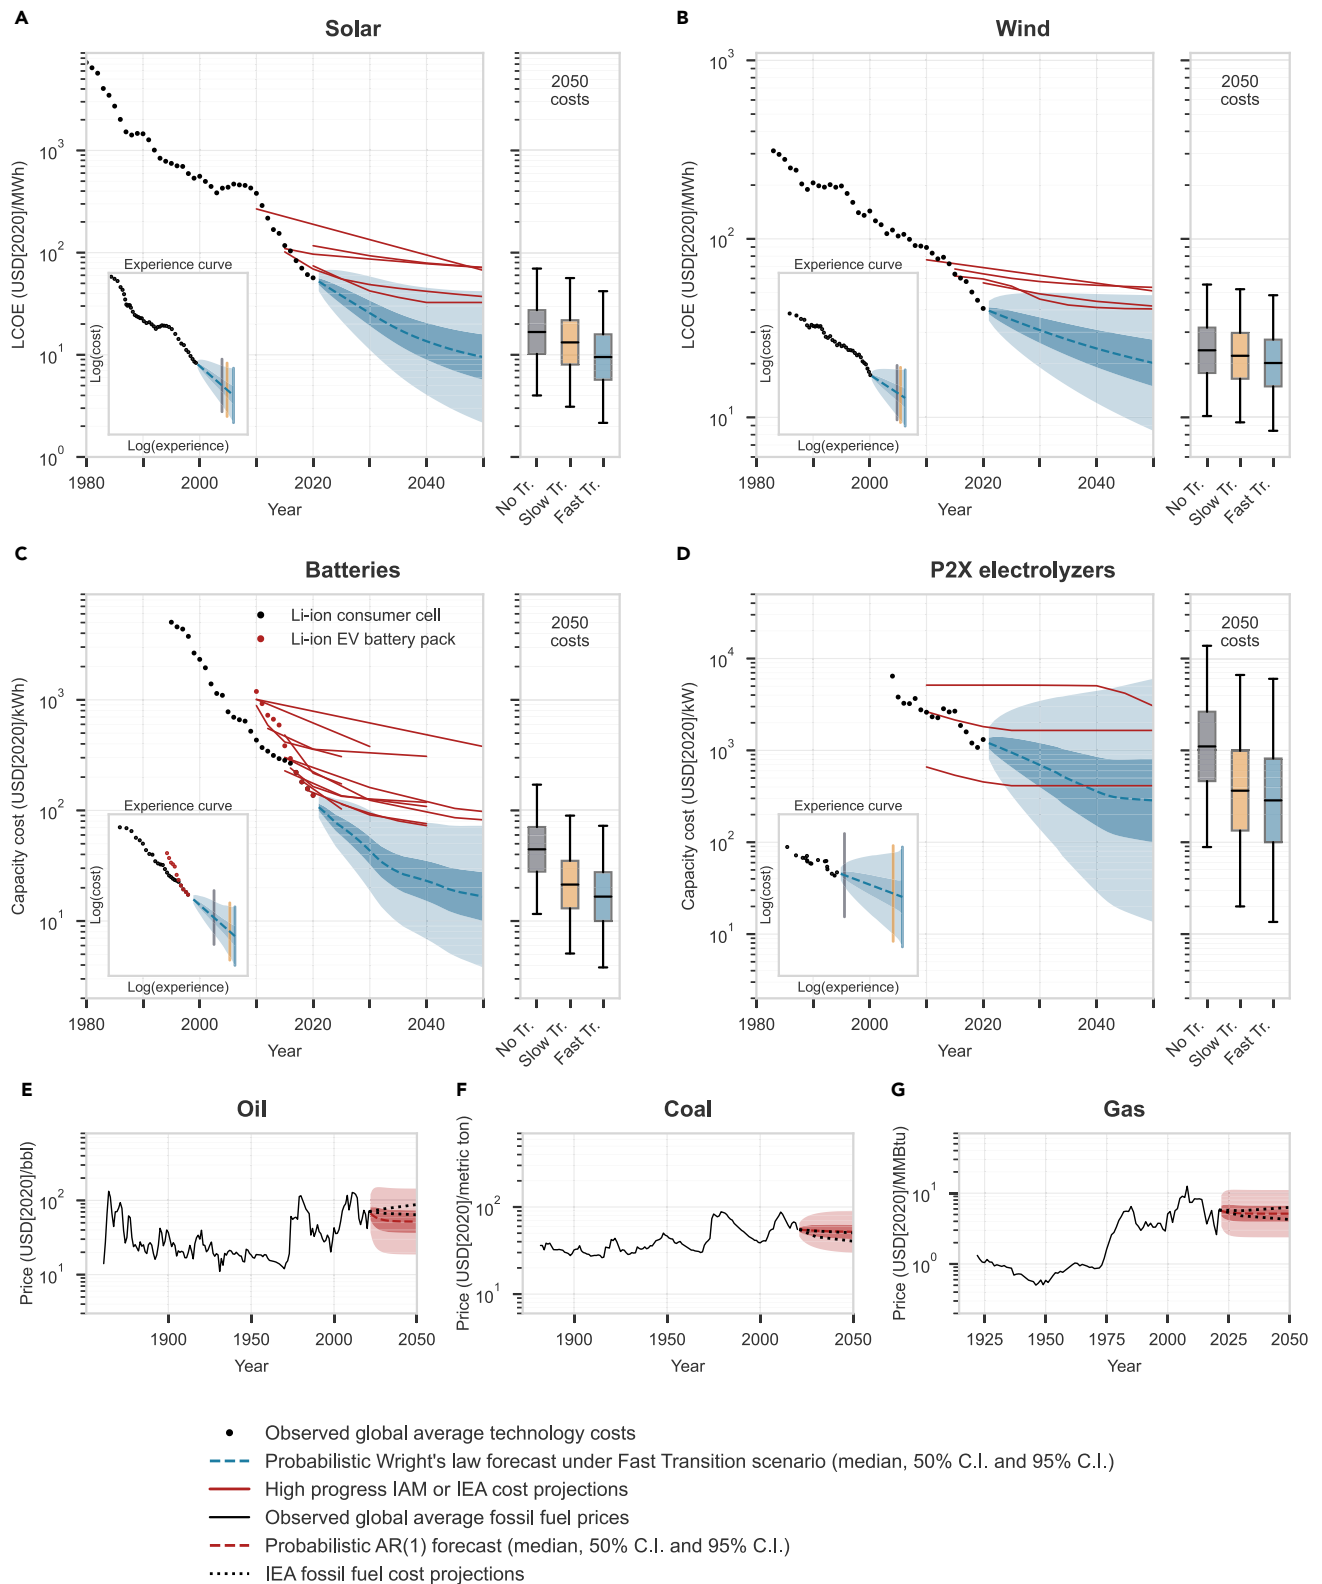


Figure 5. Technology cost forecasts

(A–D) The main panels show cost forecast distributions under the Fast Transition scenario for solar PV, wind, batteries, and PEM electrolyzers; the 50% confidence interval (CI) is dark blue, and the 95% CI is light blue. Also shown are several representative recent and past projections corresponding to

Figure 5. Continued

“optimistic” mitigation scenarios made by IAMs and the IEA (red lines) (see Figure 9). For batteries, both lithium-ion (Li-ion) consumer cells and Li-ion electric vehicle (EV) battery packs are shown, although their costs have now converged; our forecasts are based on consumer cells, whereas the IEA projections shown are based on EV batteries. The boxplots in the right-hand panels compare cost forecasts in 2050 under the No Transition, Slow Transition, and Fast Transition scenarios. The insets show historical experience curves and forecasts, with learning rates that are independent of the scenario, and vertical lines that indicate how far each technology moves down the probabilistic experience curve in each scenario. (E–G) These panels show probabilistic cost forecasts for oil, coal, and gas based on the AR(1) time-series model (see Document S1 section “Data, calibration and technology forecasts” for details of data sources and model calibration).

projects this uncertainty forward in future cost distributions. It thus provides cost ranges that are supported by empirical evidence, as opposed to the ad hoc ranges that are often used to design energy scenarios and pathways.⁴⁹ Similarly, the AR(1) model captures the historical volatility of fossil fuel prices and projects this forward in forecast distributions. Although it might appear that the range of possible outcomes in Figure 5 is larger for solar, wind, and batteries than for fossil fuels, they are actually smaller in absolute terms. The 95% confidence interval for solar costs in 2050 in the Fast Transition scenario, for example, ranges from roughly \$2 to \$40 per MWh, which is a factor of 20 and an absolute range of \$38. By contrast, the price of oil in 2050 ranges from \$20 to \$110 per barrel, which is a factor of 5.5 but an absolute range of \$90. The uncertainty ranges we forecast for fossil fuels are in line with IEA estimates.⁸¹ Note that the uncertainties for electrolyzers are much higher than for the other three key green technologies because the historical series is short and volatile.

From single technologies to a full system model

To forecast the likely costs of the green energy transition and explore how uncertainty in individual technology costs propagates through to uncertainty in system costs, we constructed a simple, transparent model of the global energy system based on well-defined technology deployment scenarios. We added seven more technologies to the seven technologies already presented: coal-fired and gas-fired electricity, nuclear power, hydroelectric power, biopower, redox flow batteries, and electricity networks. While the real-world energy system includes many other technologies, we used this limited ensemble because (1) it covers most of the current final and useful energy of the system (around 90%), (2) it includes sufficiently many diverse technologies for representing a wide range of energy transition pathways, and (3) it maintains a level of simplicity suitable for conveying our main results on future technology costs and their uncertainties.

Since our study is not intended to be comprehensive, but rather to focus on cost declines for key green technologies, we do not consider liquid biofuels, geothermal power, marine energy, traditional biomass, co-generation of heat, solar thermal energy, or CCS (our results are nevertheless robust to these modeling choices; see [Experimental procedures](#)).

Our approach to scenario construction differs from that currently used in most standard energy-economy models, where deployment in one period is used to project costs in the next, and vice versa. By iterating between costs and deployment in this way, small errors can quickly get amplified, leading to scenarios that are inconsistent with empirically observed trends. Instead, we followed earlier energy system models⁵⁰ and constructed scenarios exogenously by specifying how much energy or storage will be provided by each technology as a function of time, just as we did for single technology deployment trajectories earlier (Document S1 section “Scenario construction”). We classify energy services into four categories—transport, industry, buildings, and energy sector self-use (Document S1 section “Model

components”)—and assume that end-use sector demand grows at an overall rate of 2% per year (we vary this assumption in Document S1 section “Sensitivity to system growth rate”). We impose the constraint that all scenarios must reliably provide identical levels of energy services throughout the economy. This straightforward scenario construction method allowed us to match long-standing technology growth trends. This is in contrast to scenarios generated by IAMs, which typically do not match historical deployment trends for fast-progressing technologies and often fail to match wider transition dynamics too.⁵¹

The three scenarios that we introduced earlier—Fast Transition, Slow Transition, and No Transition—are shown in more detail in Figure 6. They run from 2021 to 2070 and were chosen to represent three distinctly different energy system pathways. In the Fast Transition scenario (Figures 6A, 6D, and 6G), solar, wind, and batteries continue to grow at rates that are somewhat slower than their long-term growth rates, as indicated in Figure 4, for approximately a decade. Electrolyzers are at an earlier stage of their S-curve; they behave similarly but stick to their current exponential growth rate for two decades. Following trajectories similar to standard S-curves, once these technologies become dominant, deployment slows to grow at 2% per year. Short-term storage and electrification of most transport are achieved with batteries, whereas long-duration energy storage (LDES) and all hard-to-electrify applications are served by power-to-X (P2X) fuels, i.e., by using electricity for hydrogen electrolysis and either directly using hydrogen or using it to make other fuels such as ammonia and methane as needed.⁵² This corresponds to an “electrify almost everything” scenario, with full sector coupling.⁵³ Under this scenario, as shown in Figure 6D, emissions quickly get close to zero. If non-energy sources of carbon emissions such as agriculture and land-use change are brought under control, it would likely meet the 1.5° Paris Agreement target (Document S1 section “Emission reductions and reduced climate risks”).

In the Slow Transition scenario (Figures 6B, 6E, and 6H), by contrast, current rapid deployment trends for key green technologies slow down immediately, so that fossil fuels are phased out more slowly and continue to dominate until mid-century. Finally, in the No Transition scenario (Figures 6C, 6F, and 6I), the energy system remains similar to its current form for several decades, as low-carbon energy sources grow only very slowly. This is similar to the reference or “no policy” scenario used by many IAMs. To provide some context for these three scenarios, they are compared with scenarios from the IPCC’s Sixth Assessment Report (AR6) in Document S1 section “Comparison with AR6 scenarios.” Full scenario details are shown in Document S1 section “Scenarios.”

To understand these scenarios, it is important to distinguish *final energy*, which is the energy delivered for use in sectors of the economy, from *useful energy*, which is the portion of final energy used to provide energy services, such as heat, light, and kinetic energy (Document S1 section “Energy system description”). Fossil fuels tend to have large end-use conversion losses in comparison to electricity, which means that significantly more final energy is required to obtain a given amount of useful energy. Switching to energy carriers with higher conversion efficiencies (e.g., moving to electric vehicles [EVs]) significantly reduces final energy consumption.^{13,54} In the Fast Transition scenario, eventually almost all energy services originate with electricity generated by solar PV and wind, which is used either directly, via batteries, or by making P2X fuels for later consumption. As shown by comparing Figures 6G and 6I, the Fast Transition qualitatively increases the role of electricity in the energy system.

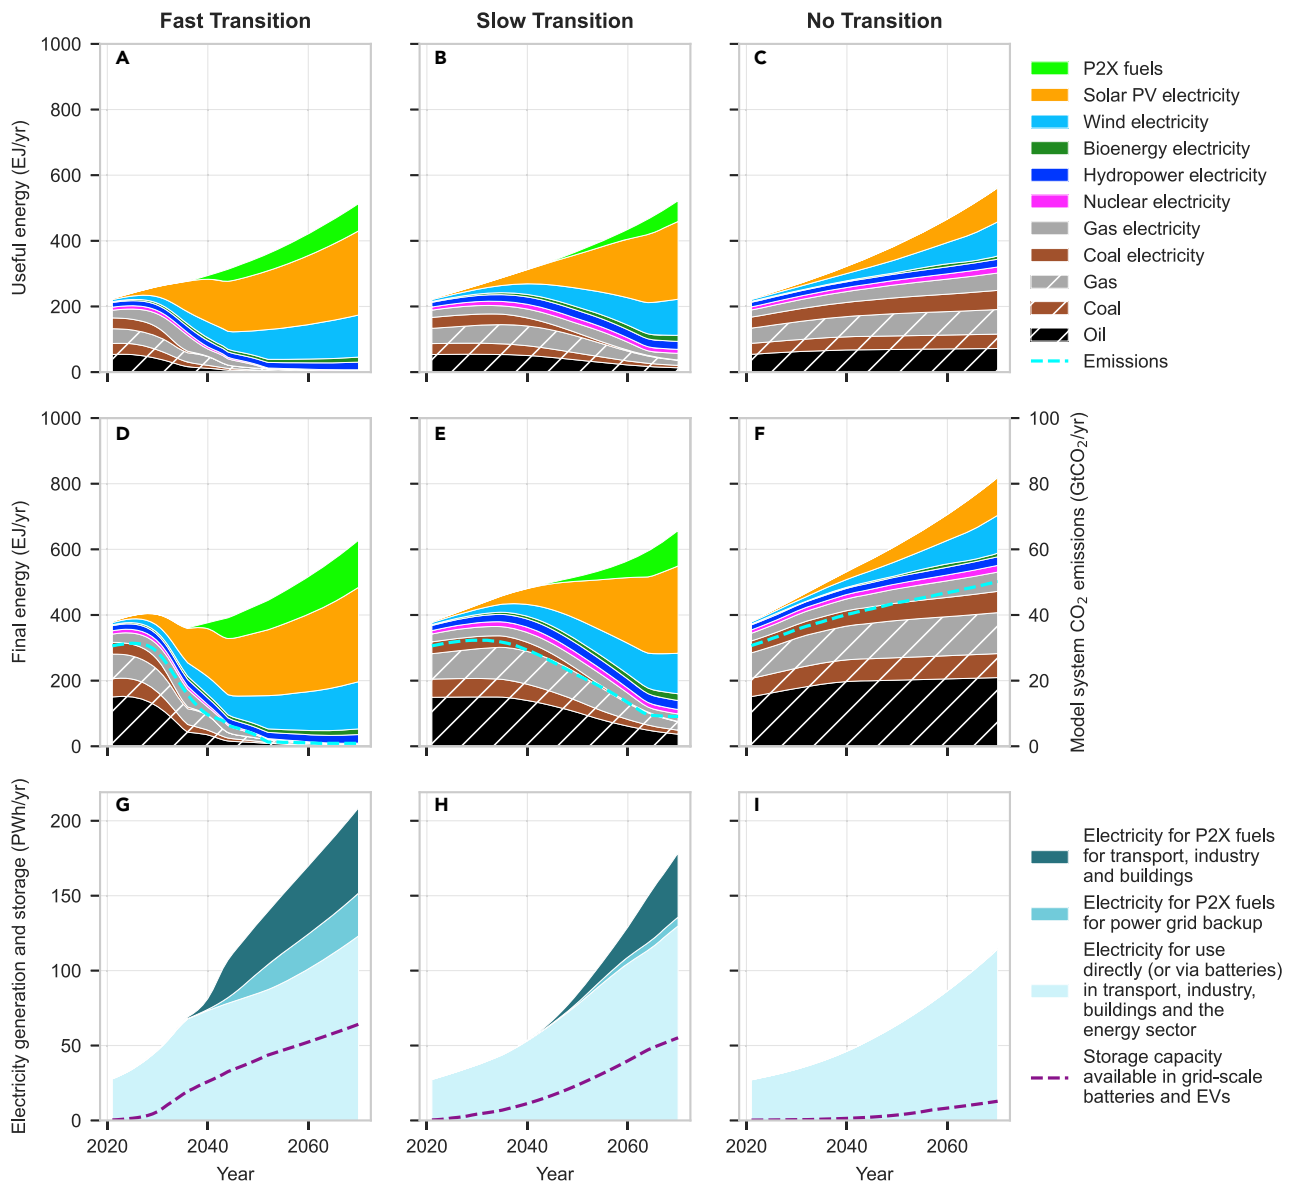


Figure 6. Scenarios

(A–I) The three columns represent the three energy system scenarios. The three rows are: (A–C) annual useful energy provided by each technology as a function of time; (D–F) annual final energy provided by each technology as a function of time; and (G–I) annual electricity generation and storage in grid-scale batteries and EV batteries. Total electricity generation is divided between final electricity delivered to the economy and electricity used to produce P2X fuels for hard-to-electrify applications and for power grid backup.

Our model captures approximately 90% of current final energy (excluding energy carriers that are already renewable, such as bioenergy and biofuels, plus petrochemical feedstock, which is not an energy carrier; see [Table S2](#)). Of course though, useful energy is what matters. The model also covers around 90% of current useful energy, but this is more difficult to estimate. Under the Fast Transition and Slow Transition scenarios, this fraction increases with time due to increased electrification. (See Document S1 section “Model description” for further details).

To estimate full system costs, we need cost forecasts for all technologies. Although coal-fired electricity and gas-fired electricity showed significant cost declines for

some of the 20th century (as power generation components underwent technological learning; see [Figure 1](#)), in the long run their costs are increasingly dominated by fuel costs,⁵⁵ so we used the AR(1) model for these (Document S1 sections “Coal electricity” and “Gas electricity”). We used stochastic Wright’s law for nuclear power, hydropower, and biopower, although these all have flat or rising costs, so the choice of cost model is immaterial here, and they have only limited significance for energy transition in this analysis. We also used stochastic Wright’s law for flow batteries, but for electricity networks, since we did not have historical cost data characterizing them in enough detail, we pessimistically assumed that unit costs will remain the same as they are now, bearing in mind that costs are heterogeneous (Document S1 section “Electricity networks”).

How much will each scenario cost?

There are many different approaches to modeling energy system pathway costs.^{56,57} We used the “direct engineering costs” approach, in which the overall cost of a scenario is computed by adding up the costs of the component technologies (Document S1 section “Estimating total system costs”). We summed the costs of direct-use oil, coal, and gas; electricity generated by seven different technologies; utility-scale grid batteries and electrolyzers; and additional infrastructure for expansion of the electricity grid. For electricity generation costs, we used the levelized cost of electricity (LCOE) metric. This is particularly advantageous here because then the experience curve formulation inherently captures historical progress trends in all LCOE components, including capital costs, capacity factors, and interest rates, which would otherwise be difficult to forecast separately (Document S1 section “Units and justification for the use of LCOE”). We estimated infrastructure costs that are not directly covered by technologies included in the model, for example, for fuel storage and distribution (Document S1 section “Fuels infrastructure”), and for fueling or charging light duty vehicles (Document S1 section “Electrification of transport”), and argue that they are roughly the same across scenarios.

To apply our probabilistic technology cost forecasting methods in a given scenario, we employed a Monte Carlo approach, simulating many different future cost trajectories, then exponentially discounting future costs to calculate the expected net present cost (NPC) of the scenario up to 2070 (Document S1 sections “Net present cost of transition” and “Main case results”). [Figure 7A](#) shows annual system costs through time for each scenario. The black boxplots represent the full cost forecast distributions, whereas the colored bars show median expenditures by technology group. This shows how, in the Fast Transition scenario, expenditures transfer rapidly from fossil fuels to key green technologies.

[Figure 7B](#) shows the annual system cost forecast distributions in 2050. Rapid replacement of fossil fuel technologies by low-cost key green technologies—in power and transport in particular—causes the expected annual energy system cost in 2050 for the Fast Transition scenario to be \$514 billion cheaper than that for the No Transition scenario, although the distribution of possible costs for the Fast Transition is wider. After 2050, as shown in [Figure 7A](#), while the median and interquartile range (IQR) remain relatively low, the uncertainty of the Fast Transition in relation to No Transition increases. If costs are in the upper end of the uncertainty range, cheaper alternatives would be used; we are not taking this into account, which is a drawback of our method.

[Figure 7C](#) shows the forecast distribution of the NPC of each scenario at a fixed discount rate of 2%. Although there is considerable uncertainty, the NPC of the Fast Transition is likely to be quite a bit lower than that of the No Transition. By contrast,

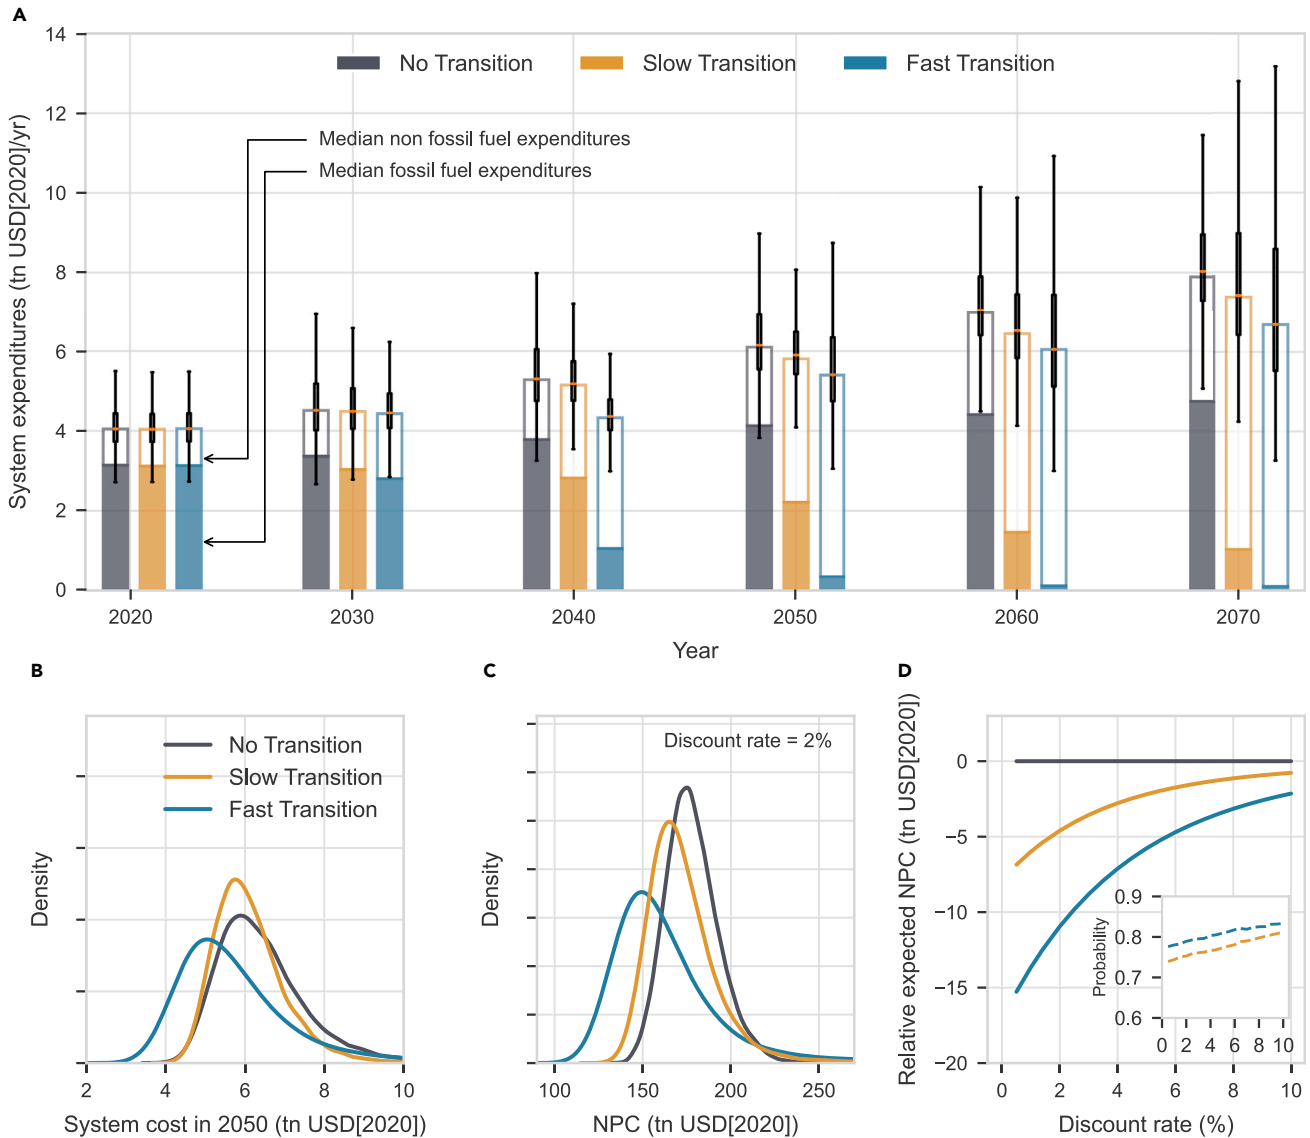


Figure 7. Scenario costs

(A) Colored bars show median annual expenditures on fossil fuel and non-fossil fuel technologies in each scenario in trillions of dollars (tn USD). Boxplots show the median and interquartile range (IQR) of total annual expenditures, and whiskers extend from the box by 1.5 times the IQR.

(B) Forecast distributions of the annual system cost in 2050 for each scenario.

(C) Forecast distributions of the net present cost (NPC) of each scenario, for a fixed discount rate of 2%.

(D) Expected net present cost of each scenario relative to the No Transition scenario, as a function of the discount rate. The inset shows the probability that the NPCs of the Fast Transition and Slow Transition will be lower than that of the No Transition, as a function of the discount rate.

the Slow Transition is not as cheap as the Fast Transition. This is because the current high spending on fossil fuels continues for decades, and the savings from key green technologies are only realized much later. Nonetheless, it also generates savings relative to the No Transition scenario. Similarly to Figure 7B, the NPC distribution of the Fast Transition is wider than that of the No Transition. Although this is caused by higher technology uncertainty, it is important to note that this increased uncertainty is compensated for by the leftward shift in the distribution, due to expected cost declines associated with scaling up key green technologies.

Figure 7D shows how the expected NPC of each scenario varies with the discount rate relative to the No Transition scenario. The inset shows that there is roughly an 80% chance that the NPC of the Fast Transition is lower than that of the No Transition, regardless of discount rate. Previous analyses have suggested that whether or not it makes good economic sense to quickly transition to clean energy technologies depends on the discount rate.^{58,59} But here we show a striking result: *the Fast Transition is likely to be much cheaper at all reasonable discount rates*. Using the 1.4% social discount rate recommended in the Stern Review,⁶⁰ for example, the expected net present saving is roughly \$12 trillion. At the higher discount rate of 5%, the expected saving is around \$5 trillion. Note that there is some evidence that technological progress does not slow when technologies reach their saturation phase.⁶¹ If this is true, then costs continue to drop at their current pace, according to Moore's law, and the Fast Transition saves even more relative to the other scenarios (see Document S1 section "Moore's law results").

We constructed an additional scenario in which nuclear plays a dominant role in replacing fossil fuels, but this is much more expensive than the other scenarios. For example, using a 1.4% discount rate, the expected NPC is about \$25 trillion more than for the No Transition (Document S1 sections "Slow nuclear transition" and "Main case results").

To enhance the credibility of our estimates, we have used consistently conservative assumptions regarding the costs, performance, and operational requirements of clean energy technologies, and we have done the opposite for fossil fuels. Our requirement that we ground forecasts on historical data means that in many cases we were forced to neglect promising solutions, such as demand-side management of power grids, heat pumps, and end-use efficiency, where there are insufficient data.¹³ As a result, the estimated savings presented here should be viewed as lower bounds on the savings likely to be achieved in reality, as many other innovative technologies and solutions are likely to be developed (Document S1 section "Estimating total system costs").

Our analysis is based on (weighted) global average costs, but there is wide geographic variation in energy costs. Within countries, solar and wind tend to be deployed first in regions where their costs are favorable, but that is not the case globally (Document S1 section "Regional differences in competing technologies"). In any case, under the Fast Transition, regional cost differences are quickly overcome through time. In the historical record of solar PV, for example, it takes less than a decade for costs to fall from the 95th to the 5th percentile of the geographical cost distribution at any fixed point in time. Because costs are summed here, global averages are sufficient to estimate total system costs, and we expect that future efforts will take advantage of geographic variation to achieve even cheaper solutions.

Although the Fast Transition happens quickly, it is still possible to replace the energy system without excessive stranding of capital. Lifetimes of large energy infrastructure projects typically range from 25 to 50 years, meaning that on average about 2%–4% of capacity needs replacing in any given year. In addition, useful energy demand grows at 2% per year in all our scenarios. These two factors make it possible for key green technologies to replace most of the existing energy system in 20 years, and replace the remaining 5% within a few decades more, without necessarily stranding assets before their economic lifetimes. Past estimates that suggest the emissions from existing, planned, and proposed electricity generation infrastructure will exceed the Paris carbon budget assumed that current utilisation rates of such assets will remain constant in future, despite an increasingly competitive market, and

that all planned and proposed deployments will go ahead, which has become increasingly unlikely over the last decade (Document S1 section “Stranded assets”).

DISCUSSION

To avoid confusion, we want to be clear about what we have done and what we have not done. In contrast to most IAMs, which attempt to project both deployment and costs conditioned on policies, we are less ambitious: we only forecast costs conditioned on deployment. Although we have tried to choose deployment scenarios that we think are reasonable, we do not attempt to forecast deployment. Our motivation for taking a less ambitious modeling approach is that this allows us to stay close to the empirical data. For improving technologies, we base all our forecasting on methods that have been carefully tested by making out-of-sample forecasts. We do not assume future technology costs, we forecast them using a well-tested methodology. If the historical data were different, the results would be different. Humility is always required in making forecasts, but we have gone to great lengths to ground our forecasts with empirical data, using statistical methods to assess their reliability.

Although our Fast Transition scenario is subjective, we believe it is plausible (Document S1 section “Is the speed of the Fast Transition achievable?”). The deployment trajectories are in line with past trends. There appear to be no major obstacles to bringing the necessary technologies to scale in terms of land use, sea, climate, raw materials, manufacturing capacity, energy return on energy invested, or system integration.⁶² Nonetheless, there are significant institutional challenges, and to stay on the current growth paths for the next decade, policies that enforce portfolio standards and/or stimulate demand will likely be needed. Our key contribution here is to show that if we can stay on these growth paths for the next decade, we will likely realize substantial savings. The cornerstone of the Fast Transition scenario is the timely expansion of key green technologies, because only as these are scaled up can fossil fuels be phased out and the savings be realized. The primary policy implication of our results is that there are enormous advantages to rapid deployment of key green technologies. Achieving this is likely to require strong international policies for building infrastructure, skills training, and making the investments required to realize future gains.

Our approach is complementary to IAMs. It builds on historical trends directly and thus provides a counterweight to projections by IAMs. We have demonstrated that the constraints that are commonly used in IAMs are likely an important cause of the mismatch of their projections with historical data. Future work could explore how softening these constraints within IAMs changes their projections.

We want to stress that, unlike IAMs, we are not attempting to find optimal solutions. There are very likely other scenarios that are cheaper than the Fast Transition scenario, which was constructed to explore whether (with sufficiently rapid deployment) a rapid transition can achieve net cost savings, and if so, with what probability. Given the likely future cost of gas, it could be possible to achieve cheaper scenarios by using gas in place of P2X fuels in some locations and applications, but these of course would not be zero-emissions systems. Similarly, while fossil fuel prices have not historically trended down, competition from key green technologies may force them down, although this is feasible only at substantially reduced production levels where only the cheapest fossil fuel producers are competitive.⁶³ This emphasizes the point that, while most of the Fast Transition is aligned with market forces, policies that discourage the use of fossil fuels will still be needed to fully decarbonize energy.

Our forecasts are based on time-series methods. It would be preferable to be able to make forecasts based on first principles, but unfortunately this is not possible (even though Wright's law can be derived from a simple model⁶⁴). Certain characteristics of technologies, such as modularity, are predicted to be associated with more rapid technological progress;⁶⁴ but despite some evidence that this is indeed true for modularity, it only explains a small fraction of the variance.⁶⁵ A fundamental limitation of all model-based forecasting methods is data, as this is required for model calibration. When historical data are too sparse or do not exist, these methods can not be used. In such cases, expert elicitation methods may be applied,⁶⁶ although this requires care, as retrospective analysis indicates that time-series methods are more reliable.⁶⁷ Future research to address the issue of sparse data could consider hybrid approaches that combine model-based forecasts with expert judgments, or it could investigate the extent to which technology analogs can reliably be used in forecasting. Perhaps most useful though would be a better understanding of the role of technology aggregation in model-based forecasting; for example, how do forecasts of global solar PV costs relate to specific subtypes of PV and to regional costs? In any case, the time-series methods that we used here are currently the most reliable way to make conditional forecasts of technology costs that agree with historical data.

Although the infrastructure costs for a rapid green energy transition are substantial, we forecast that they are likely to be more than offset by lower energy costs. The largest infrastructure cost is for enhanced grid capacity. In 2050, for example, our estimated electricity network annual expenditure for the Fast Transition is about \$670 billion per year, compared with \$530 billion per year for the No Transition. However, the expected total system cost in 2050 is about \$5.9 trillion per year for the Fast Transition and \$6.3 trillion per year for the No Transition. Thus, although the additional \$140 billion of grid costs might seem expensive, it is significantly less than the savings due to cheaper energy. The essential reason that the Fast Transition is cheaper than the Slow Transition is because it realizes the cost savings due to cheaper energy sooner—faster deployment increases the probability of rapid progress in key green technologies, so that savings accrue for longer.

The likelihood of cheaper energy raises the possibility of a rebound effect. Cheaper energy may increase global energy demand so that it grows faster than the historical 2% per year rate assumed here. We view this as a "good problem": while this would raise overall costs in the Fast Transition, renewables produce clean energy, and cheaper energy is likely to improve global living standards. To address the potential cost increases, we performed a sensitivity analysis around the system growth rate assumption, by also considering long-run growth rates of 1% and 3%. The resulting energy systems vary widely in size and therefore represent a wide range of plausible population, economic, and technological pathways in future. We found that our results are robust to these variations (Document S1 section "Sensitivity to system growth rate").

In response to our opening question, "Is there a path forward that can get us to net-zero emissions cheaply and quickly?," our answer is: "Very likely, and the savings are probably quite large." Our quantitative analysis supports other recent efforts using up-to-date data and technology assumptions that conclude that the green energy transition may be cheap.^{16–19,68–71} The 2022 IPCC AR6 estimates that the additional cost of decarbonizing the energy system in order to have a greater than 67% chance of keeping warming below 2°C corresponds to a GDP loss in 2050 of 1.3%–2.7%.⁴⁰ Our results suggest that there is likely no cost at all—the transition is expected to be a net economic benefit, raising future GDP.

We have demonstrated that the models used by the IPCC have in the past consistently overestimated key green technology costs. It is important that this bias is addressed (although, as shown in Figure 9A, the PV cost projections in the AR6 database have so far exhibited the same upward bias seen previously). In the context of the probabilistic forecasts presented in this paper, the IPCC database models are only considering costs that are extremely unlikely—in the pessimistic direction—and are not fully exploring the space of plausible scenarios. IPCC conclusions thus appear to be based on an over-sampling of near worst-case scenarios regarding key green technology costs. While even technologies with strong progress trends sometimes experience cost increases, as has occurred for solar in the mid-2000s and the current period due to supply shortages of key production inputs, our results account for these fluctuations. By carefully characterizing historical progress trends and volatility, the stochastic methods used here capture both the downside risk that progress in some key green technologies may stall, and the upside risk—the probability that via routine invention and innovation some technology costs will fall faster than historical trends some of the time.

Our analysis indicates that even the downside outcomes of a rapid green energy transition are not that bad, due to the dramatic cost declines seen already. When energy system pathways are viewed in terms of bets placed on portfolios of technologies,⁷² we find that the Fast Transition scenario has an expected payoff of around \$5–\$15 trillion. Moreover, it is also a safe bet, with around an 80% probability that it will be cheaper than continuing with a fossil fuel-based system (and 82% when compared with a slower transition). Since the future is uncertain, all public policy and decision-making is ultimately a question of making the smartest bets we can, given the often precarious circumstances we face. Our results suggest that deploying technologies according to the Fast Transition scenario is a very good bet, both in terms of lowest costs and lowest emissions.

We want to emphasize that our results indicate that a rapid green energy transition is likely to be beneficial, even if climate change were not a problem. When climate change is taken into account, the benefits of the Fast Transition become overwhelming. A common simplified method for estimating economic damages due to climate change is to apply a social cost of carbon (SCC)^{15,73–75} to emissions. The range of proposed values is vast, but just as an example, at a discount rate of 5%, assuming SCC values in the range \$30–300/tCO₂ (rising at 3% per year⁵⁸) yields total expected Fast Transition savings, up to 2070, of \$31–\$255 trillion. At a lower discount rate of 1.4%, the range of expected savings is \$88–\$775 trillion. Thus, the benefits of the Fast Transition are likely much larger than the energy system cost savings evaluated in this study.

The belief that the green energy transition will be expensive has been a major driver of the ineffective response to climate change for the past 40 years. This pessimism is at odds with past technological cost improvement trends and risks locking humanity into an expensive and dangerous energy future. While arguments for a rapid green transition cite benefits such as the avoidance of climate damages, reduced air pollution, and lower energy price volatility (Document S1 section “Additional benefits from the Fast Transition”), these benefits are often contrasted against discussions about the associated costs of the transition. Our analysis suggests that such trade-offs are unlikely to exist: *a greener, healthier, and safer global energy system is also likely to be cheaper*. Updating expectations to better align with historical evidence could fundamentally change the debate

about climate policy and dramatically accelerate progress to decarbonize energy systems around the world.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Correspondence and requests for resources should be addressed to rupert.way@smithschool.ox.ac.uk.

Materials availability

This study did not generate any new materials.

Time-series models

We employ two time-series models for forecasting technology costs. The first is a first-difference stochastic form of *Wright's law*, developed and tested by Lafond et al.,³⁹ which models costs dropping as a power law of cumulative production. Let c_t be the cost and z_t be the experience of a given technology at time t , and let $u_t \sim \mathcal{N}(0, \sigma_u^2)$ be an independent and identically distributed (IID) draw from a normal distribution. Then future costs are predicted using the iterative relationship

$$\log c_t - \log c_{t-1} = -\omega(\log z_t - \log z_{t-1}) + u_t + \rho u_{t-1}. \quad (\text{Equation 1})$$

This relationship has three parameters. For a given technology, the experience exponent ω characterizes the average rate at which costs drop as a function of experience, and the noise variance σ_u^2 characterizes the variability of this relationship. The autocorrelation parameter ρ characterizes the persistence of fluctuations in cost improvements. To avoid overfitting, and to ensure that our forecasts adhere strictly to the same statistical properties as those tested by Lafond et al.³⁹ we use $\rho = 0.19$ for all technologies, which was found to be a good overall choice for 50 different technologies. This was necessary because fitting three parameters to short data series such as those we have here degrades out-of-sample forecasting accuracy. (We also did a comparison of all our results replacing Wright's law by a generalized form of Moore's law [see Document S1 section "Moore's law results"]).

When applying the model to technologies with falling costs, as shown in [Figure 5](#), two features of the model must be stressed. First, the Wright's law model does not simply "assume" that if costs fell in the past then they will fall in future—indeed, costs are predicted to rise with a non-zero probability that depends directly on observed data in the past. Second, despite the downward trends, all cost forecast distributions are always strictly positive, since costs develop in log space.

For fossil fuels we use an AR(1) model:

$$\log c_t = \log c_{t-1} + \beta(\mu - \log c_{t-1}) + \epsilon_t, \quad \text{with IID } \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \quad (\text{Equation 2})$$

where $\mu = \mathbb{E}[\log c_t]$ is the unconditional mean of the logarithm of cost, σ_ϵ is the volatility of the noise process ϵ_t , and β is the rate of mean reversion. For comprehensive details on forecasting methods see Document S1 section "Technology cost models."

Additional technology cost projections

[Figures 8](#) and [9](#) show the solar and wind cost projection ensembles that underlie the specific cost projections displayed in [Figures 3](#) and [5](#).

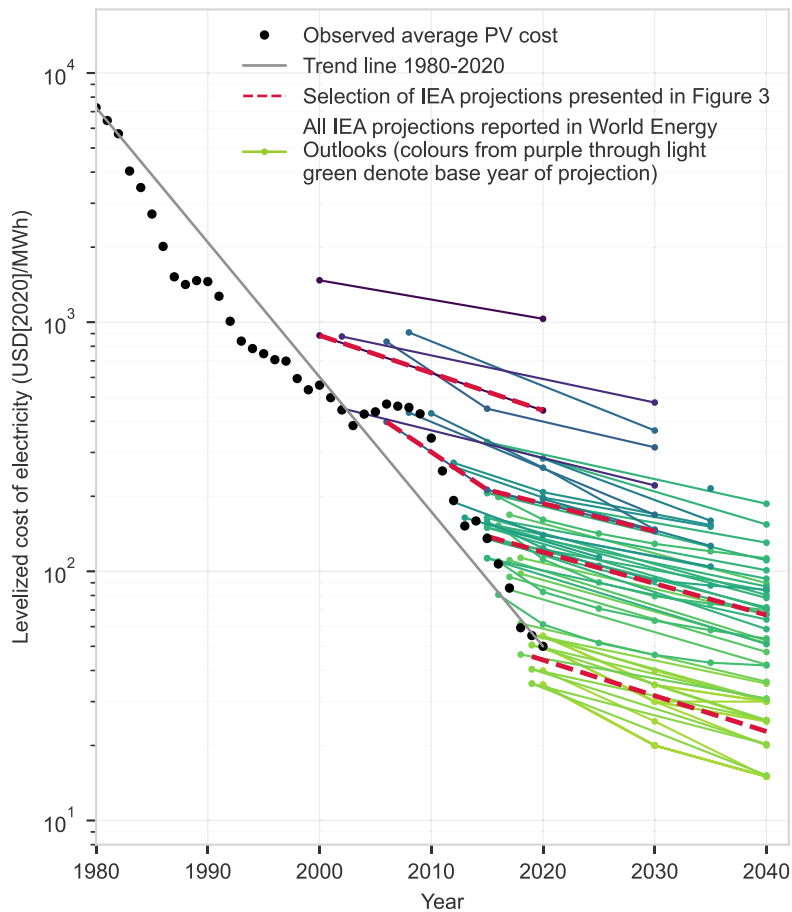


Figure 8. IEA PV LCOE projections

All PV LCOE projections found in the IEA's World Energy Outlook (WEO) reports are shown in colors varying from purple through light green (note that "projection" here means conditional forecast—this is a forecast that is conditional upon a whole array of modeling assumptions regarding the scenario within which the forecast is embedded). The first such projection was found in the WEO 2001. The four projections we selected to plot in Figure 3 are shown in red and were chosen as examples of "high progress" projections. The first two, published in the WEOs from 2001 and 2008, may be considered high progress projections, because in those reports, cost ranges were provided, and we simply picked the lowest points of those ranges. The upper ends of the ranges were significantly higher. The second two (beginning in 2015 and 2019) may be interpreted as "high progress" projections, because they correspond to the highest mitigation scenarios available in the WEOs from which they are sourced (WEO 2016 and 2020). Note, however, that in those reports, only region-specific cost projections were provided, so we have plotted the simple global average of those values in the high mitigation scenarios. Observed values are from the Performance Curve Database (described in Nagy et al.²⁵) up to 2010 and from Bloomberg New Energy Finance (BNEF) thereafter.

See Document S1 section "Data, calibration, and technology forecasts" for more details on data sources.

Scenario construction

We model the supply of energy services increasing at a fixed rate per year (2% in our main specification). Energy services include heating/cooling, light, mobility, sustenance, materials, hygiene, and communications, but since these are hard to measure and data are sparse, we take useful energy as a proxy for energy services. Energy transition scenarios are constructed by assuming that growth rates of energy carriers and technologies follow logistic ("S") curves with specified start and end points consistent with the growth of total useful energy. We model the relationship

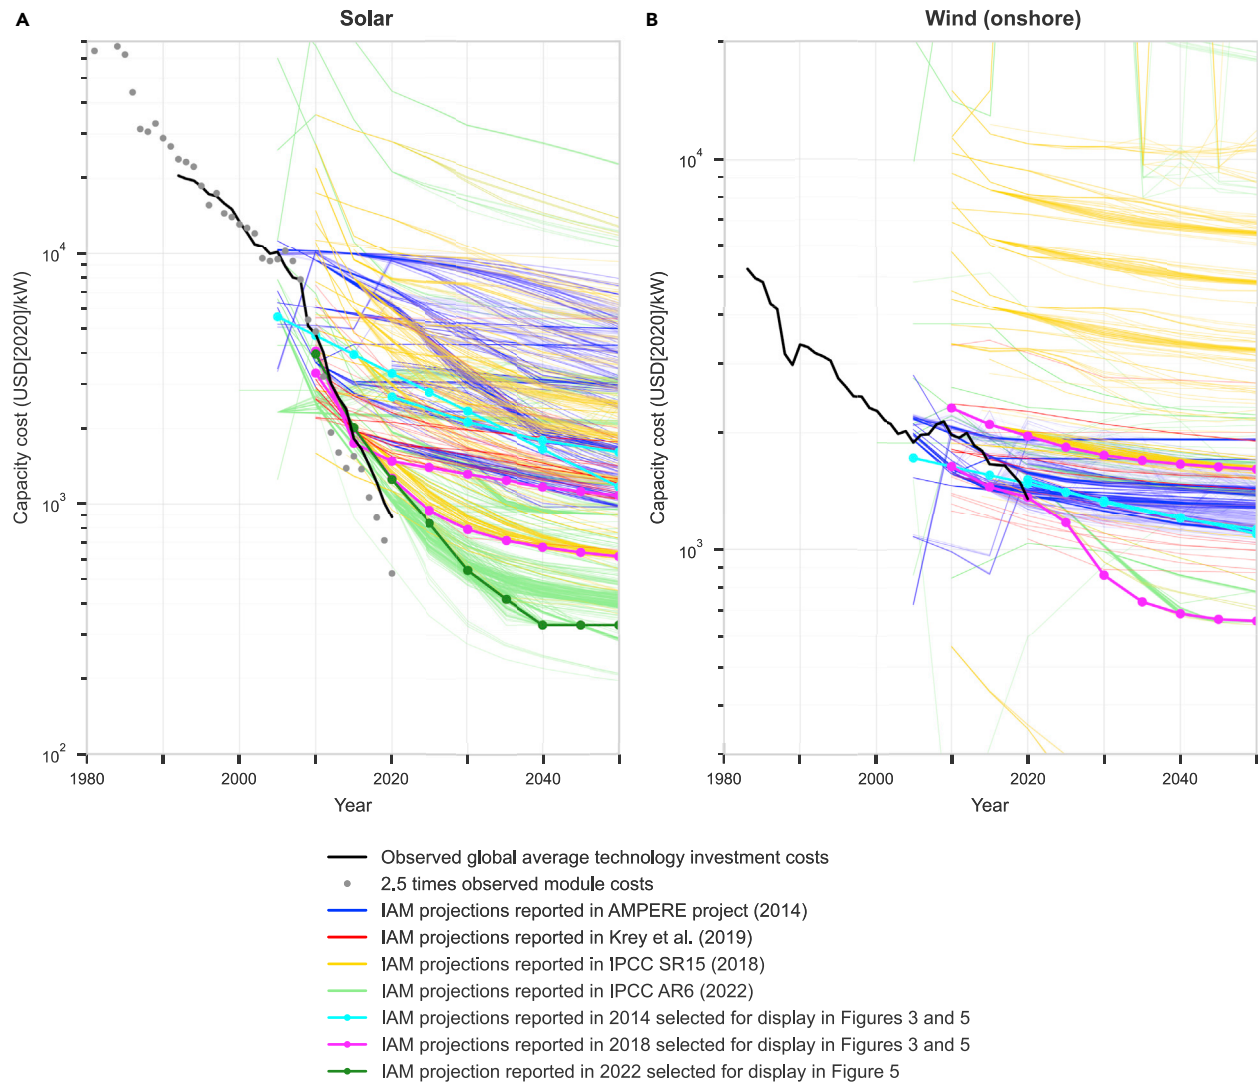


Figure 9. PV and wind capital cost projections reported by IAMs

Capital cost projections reported by various modeling comparison projects are shown as blue, red, yellow, and green lines for (A) PV and (B) onshore wind. Each line corresponds to a single scenario. To construct and plot the LCOE projections in Figure 3, we selected two capacity cost projections reported in 2014 (cyan lines) and two reported in 2018 (magenta lines). For Figure 5, we also added one PV projection reported in 2022 (dark green). These may all be interpreted as “high progress” projections because they are among the lowest in their cohorts (i.e., the cyan lines are on the low end of the suite of blue lines, the magenta lines are low relative to the red and yellow lines, and the dark green line is low relative to the green lines). Note that these are (in eight out of nine cases) global average values, whereas some other projections are region specific. For PV, the projections plotted are as follows: (1) model: MESSAGE, scenario: “AMPERE3-450,” region: World (from AMPERE⁴¹); (2) model: DNE21, scenario: “AMPERE3-450,” region: World (from AMPERE⁴¹); (3) model: IMAGE 3.0, scenario: baseline, region: China (from Krey et al.⁴²); (4) model: REMIND-MAGPIE 1.7–3.0, scenario: SMP_1p5C_early, region: World (from SR15⁴³); and (5) model: WITCH 5.0, scenario: EN_INDCi2030_1000, region: World (from AR6⁴⁰). For wind, the projections plotted are as follows: (1) model: MESSAGE, scenario: AMPERE3-450, region: World (from AMPERE⁴¹); (2) model: DNE21, scenario: AMPERE3-450, region: World (from AMPERE⁴¹); (3) model: AIM/CGE 2.1, scenario: TERL_15D_LowCarbonTransportPolicy, region: World (from SR15⁴³); and (4) model: REMIND-MAGPIE 1.7–3.0, scenario: SMP_1p5C_early, region: World (from SR15⁴³). To calculate LCOEs, we used technology lifetimes and discount rates reported in Krey et al.,⁴² and operations and maintenance (O&M) values from the original studies where possible (and if not, then Krey et al.⁴² values again). We used global average capacity factors of 0.18 for PV and 0.3 for wind, based on recent data reported by IRENA⁸³ and IEA.⁷⁶ Observed data sources for PV are given in Table S21. Wind data are from IRENA.⁸³

between final and useful energy based on average conversion efficiency factors given by DeStercke.⁷⁷ We take these conversion factors to be static and apply them on a per-energy-carrier, per-sector basis. The model therefore does not

include improvements in the underlying efficiency of each energy carrier at providing energy services, but it does allow efficiency improvements on a per-sector basis, via energy carrier substitution. For example, the conversion efficiencies of oil and electricity in the transport sector are assumed constant, yet the efficiency of the sector as a whole may be improved by switching from oil to electricity. The end point of each scenario in 2070 is defined by the shares of technologies providing electricity generation and the shares of energy carriers providing final energy. The start points for all scenarios are identical and match the current shares in the base year, 2020. Details of all growth rates, timings, and energy carrier mixes for each scenario are given in Document S1 sections “Scenario construction” and “Scenarios.”

Our modeling approach is based on two key design principles: (1) include only the minimal set of variables necessary to represent most of the global energy system and the most important cost and production dynamics, and (2) ensure all assumptions and dynamics are technically realistic and closely tied to empirical evidence (Document S1 section “General approach”). This means that we focus on energy technologies that have been in commercial use for sufficient time to develop a reliable historical record for forecasting purposes. This is also an essential constraint because deployment of technologies typically takes a long time, and only technologies that have track records are positioned to play an immediate role in confronting climate change.²²

We choose a level of model granularity well suited to the probabilistic forecasting methods used, i.e., one that allows accurate model calibration and ensures overall cost-reduction trends associated with cumulative production are captured for each technology. Our model design can be run on a laptop, is easy to understand and interpret, and allows us to calibrate all components against historical data so that the model is firmly empirically grounded. The historical data do not exist to do this on a more granular level.

We omitted several minor energy technologies. Co-generation of heat, traditional biomass, marine energy, solar thermal energy, and geothermal energy were omitted either due to insufficient historical data or because they have not exhibited significant historical cost improvements, or both. Liquid biofuels were also excluded because any significant expansion would have high environmental costs (Document S1 section “Bioenergy, solar thermal energy, marine energy and geothermal energy”). Finally, CCS in conjunction with fossil fuels was omitted because (1) it is currently a very small, low growth sector, (2) it has exhibited no promising cost improvements so far in its 50-year history, and (3) the cost of fossil fuels provides a hard lower bound on the cost of providing energy via fossil fuels with CCS (Document S1 section “CCS”). This means that within a few decades electricity produced with CCS will likely not be competitive even if CCS is free. There may of course be some role for CCS in non-energy, direct-emission applications, but this is outside the scope of this paper.

Since renewable energy production is variable, storage is essential. In the Fast Transition scenario we have allocated so much storage capacity using batteries and P2X fuels that the entire global energy system could run for a month without any sun or wind (Document S1 section “Energy storage and flexibility requirements”). This is a sensible choice because both batteries and electrolyzers have highly favorable trends for cost and production (Document S1 sections “Batteries” and “Hydrogen and electrolyzers”). From 1995 to 2018 the production of lithium-ion (Li-ion) batteries increased at 30% per year, while costs dropped at 12% per year, giving an

experience curve comparable to that of solar PV.⁷⁸ Currently, about 60% of the cost of electrolytic hydrogen is electricity, and hydrogen is around 80% of the cost of ammonia,⁷⁹ so these automatically take advantage of the high progress rates for solar PV and wind.

We ensure system reliability constraints are met—including robustness to seasonal demand variations—by providing sufficient levels of energy storage, firm capacity resources, over-generation of variable renewable energy (VRE) sources, and network expansion⁸⁰ (Document S1 section “Energy storage and flexibility requirements”). To be specific, when VRE penetration is high, we ensure enough utility-scale battery storage is available to store 20% of average daily electricity generation (though note that daily generation is much higher than daily end-use consumption, because excess generation is used to produce P2X fuels). Flow batteries are able to store a further 10% of average daily generation. In addition, when VRE penetration is high, transport is electrified, which as well as being a flexible demand source, could also act as another storage source (though system reliability constraints are met here without relying on it). Excess VRE is used to produce P2X fuels in sufficient quantities to supply all end-use sector requirements and also to provide global power grid backup for 1 month each year.

Data and code availability

We used data from a wide range of sources. Many of these were free and openly available on the internet, but some were accessed via standard university-wide subscription licenses held by the University of Oxford. Sources include: the IEA,⁸¹ BP,⁸² the International Renewable Energy Agency,⁸³ Lazard,⁸⁴ the U.S. Energy Information Administration,⁸⁵ Bloomberg New Energy Finance, Bloomberg L.P. (via Bloomberg Terminal), and several academic papers. For more details, see Document S1 section “Data, calibration and technology forecasts.” All data will be made available upon request (unless legal restrictions exist).

The code used in this analysis will be made available upon request.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.joule.2022.08.009>.

ACKNOWLEDGMENTS

This work was supported by funding from Partners for a New Economy (R.W. and J.D.F.); Baillie Gifford (J.D.F.); the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 730427 (COP21 RIPPLES) (R.W.); and the Oxford Martin School, through the Institute of New Economic Thinking and the Post-Carbon Transition programme (LDR00530) (M.C.I. and P.M.). This research was also funded by the Economics of Energy Innovation and System Transition project (EEIST), which is jointly funded through UK Aid by the UK Government Department for Business, Energy, and Industrial Strategy (BEIS) and the Children’s Investment Fund Foundation (CIFF) (R.W. and M.C.I.). The contents of this manuscript represent the views of the authors and should not be taken to represent the views of the UK government or CIFF. The authors gratefully acknowledge all these sources of financial support and, additionally, the Institute for New Economic Thinking at the Oxford Martin School for its continuing support. The authors also thank Lucas Kruitwagen and Jing Meng for valuable contributions to the work.

AUTHOR CONTRIBUTIONS

Conceptualization, J.D.F.; methodology, R.W. and J.D.F.; software, R.W. and M.C.I.; investigation, R.W.; data curation, R.W. and M.C.I.; formal analysis, R.W. and M.C.I.; writing—original draft, J.D.F., R.W., M.C.I., and P.M.; writing—review & editing, J.D.F., R.W., M.C.I., and P.M.; visualization, R.W.; supervision, J.D.F.; funding acquisition, J.D.F.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: December 23, 2021

Revised: May 16, 2022

Accepted: August 19, 2022

Published: September 13, 2022

REFERENCES

- Perlin, J. (1999). *From Space to Earth: the Story of Solar Electricity* (Aatech Publications).
- Creutzig, F., Agoston, P., Goldschmidt, J.C., Luderer, G., Nemet, G., and Pietzcker, R.C. (2017). The underestimated potential of solar energy to mitigate climate change. *Nat. Energy* 2, 17140. <https://doi.org/10.1038/nenergy.2017.140>.
- Xiao, M., Junne, T., Haas, J., and Klein, M. (2021). Plummeting costs of renewables - are energy scenarios lagging? *Energy Strategy Rev* 35, 100636. <https://doi.org/10.1016/j.esr.2021.100636>.
- Jaxa-Rozen, M., and Trutnevyte, E. (2021). Sources of uncertainty in long-term global scenarios of solar photovoltaic technology. *Nat. Clim. Change* 11, 266–273. <https://doi.org/10.1038/s41558-021-00998-8>.
- Hoekstra, Auke, Steinbuch, M., and Verbong, G. (2017). Creating agent-based energy transition management models that can uncover profitable pathways to climate change mitigation. *Complexity* 2017, 1–23. <https://doi.org/10.1155/2017/1967645>.
- Shiraki, H., and Sugiyama, M. (2020). Back to the basic: toward improvement of technoeconomic representation in integrated assessment models. *Clim. Change* 162, 13–24. <https://doi.org/10.1007/s10584-020-02731-4>.
- Victoria, M., Haegel, N., Peters, I.M., Sinton, R., Jäger-Waldau, A., del Cañizo, C., Breyer, C., Stocks, M., Blakers, A., Kaizuka, I., et al. (2021). Solar photovoltaics is ready to power a sustainable future. *Joule* 5, 1041–1056. <https://doi.org/10.1016/j.joule.2021.03.005>.
- Stern, N. (2016). Economics: current climate models are grossly misleading. *Nature* 530, 407–409. <https://doi.org/10.1038/530407a>.
- Pindyck, R.S. (2013). Climate change policy: what do the models tell us? *J. Econ. Lit.* 51, 860–872. <https://doi.org/10.1257/jel.51.3.860>.
- Farmer, J.D., Hepburn, C., Mealy, P., and Teytelboym, A. (2015). A third wave in the economics of climate change. *Environ. Resour. Econ.* 62, 329–357.
- Gambhir, A., Butnar, I., Li, P.-H., Smith, P., and Strachan, N. (2019). A review of criticisms of integrated assessment models and proposed approaches to address these, through the lens of beccs. *Energies* 12, 1747. <https://doi.org/10.3390/en12091747>.
- McCollum, D.L., Gambhir, A., Rogelj, J., and Wilson, C. (2020). Energy modellers should explore extremes more systematically in scenarios. *Nat. Energy* 5, 104–107. <https://doi.org/10.1038/s41560-020-0555-3>.
- Lovins, A.B., Ürge-Vorsatz, D., Mundaca, L., Kammen, D.M., and Glassman, J.W. (2019). Recalibrating climate prospects. *Environ. Res. Lett.* 14, 120201.
- Pye, S., Broad, O., Bataille, C., Brockway, P., Daly, H.E., Freeman, R., Gambhir, A., Geden, O., Rogan, F., Sanghvi, S., et al. (2021). Modelling net-zero emissions energy systems requires a change in approach. *Clim. Policy* 21, 222–231. <https://doi.org/10.1080/14693062.2020.1824891>.
- Stern, N., Stiglitz, J.E., and Taylor, C. (2021). The economics of immense risk, urgent action and radical change: towards new approaches to the economics of climate change. *J. Econ. Methodol.* 29, 1–36.
- Grubler, A., Wilson, C., Bento, N., Boza-Kiss, B., Krey, V., McCollum, D.L., Rao, N.D., Riahi, K., Rogelj, J., De Stercke, S., et al. (2018). A low energy demand scenario for meeting the 1.5°C target and sustainable development goals without negative emission technologies. *Nat. Energy* 3, 515–527. <https://doi.org/10.1038/s41560-018-0172-6>.
- Victoria, M., Zhu, K., Brown, T., Andresen, G.B., and Greiner, M. (2020). Early decarbonisation of the European energy system pays off. *Nat. Commun.* 11, 6223. <https://doi.org/10.1038/s41467-020-20015-4>.
- He, G., Lin, J., Sifuentes, F., Liu, X., Abhyankar, N., and Phadke, A. (2020). Rapid cost decrease of renewables and storage accelerates the decarbonization of china's power system. *Nat. Commun.* 11, 2486. <https://doi.org/10.1038/s41467-020-16184-x>.
- Bogdanov, D., Ram, M., Aghahosseini, A., Gulagi, A., Oyewo, A.S., Child, M., Caldera, U., Sadovskaia, K., Farfan, J., De Souza Noel Simas Barbosa, L., et al. (2021). Low-cost renewable electricity as the key driver of the global energy transition towards sustainability. *Energy* 227, 120467. <https://doi.org/10.1016/j.energy.2021.120467>.
- Arthur, W.B. (2011). *The Nature of Technology. What It Is and How It Evolves* (Free Press).
- Fouquet, R. (2008). *Heat, Power and Light: Revolutions in Energy Services* (Edward Elgar Publications).
- Grubler, A., Nakićenović, N., and Victor, D.G. (1999). Dynamics of energy technologies and global change. *Energy Policy* 27, 247–280. [https://doi.org/10.1016/S0301-4215\(98\)00067-6](https://doi.org/10.1016/S0301-4215(98)00067-6).
- Koh, Heebyoung, and Magee, C.L. (2006). A functional approach for studying technological progress: application to information technology. *Technol. Forecasting Soc. Change* 73, 1061–1083. <https://doi.org/10.1016/j.techfore.2006.06.001>.
- Koh, H., and Magee, C.L. (2008). A functional approach for studying technological progress: extension to energy technology. *Technol. Forecasting Soc. Change* 75, 735–758. <https://doi.org/10.1016/j.techfore.2007.05.007>.
- Nagy, B., Farmer, J.D., Bui, Q.M., and Trancik, J.E. (2013). Statistical basis for predicting technological progress. *PLoS One* 8, e52669.
- U.S. Geological Survey (2022). Mineral Commodity Summaries 2022. Technical report. <http://pubs.er.usgs.gov/publication/mcs2022>.
- Trancik, J.E., Jean, J., Kavlak, G., Klemun, M.M., Edwards, M.R., McNerney, J., Miotti, M., Brown, P.R., Mueller, J.M., and Needell, Z.A. (2015). Technology Improvement and Emissions Reductions as Mutually Reinforcing Efforts: Observations from the Global Development of Solar and Wind Energy. Working paper (MIT Press).
- Reiner, D.M. (2016). Learning through a portfolio of carbon capture and storage demonstration projects. *Nat. Energy* 1, 1–7.

29. Rubin, E.S., Davison, J.E., and Herzog, H.J. (2015). The cost of CO₂ capture and storage. *Int. J. Greenhouse Gas Control* 40, 378–400. Special Issue commemorating the 10th year anniversary of the publication of the Intergovernmental Panel on Climate Change Special Report on CO₂ Capture and Storage. <https://doi.org/10.1016/j.ijggc.2015.05.018>.
30. Malkiel, B.G. (1973). *A Random Walk down Wall Street* (Norton & Co).
31. Pindyck, R.S. (1999). The long-run evolution of energy prices. *Energy J* 20, 19449089. <http://www.jstor.org/stable/41322828>.
32. Shafiee, S., and Topal, E. (2010). A long-term view of worldwide fossil fuel prices. *Appl. Energy* 87, 988–1000. <https://doi.org/10.1016/j.apenergy.2009.09.012>.
33. Farmer, J.D., and Lafond, F. (2016). How predictable is technological progress? *Res. Policy* 45, 647–665.
34. Moore, G.E. (1998). Cramming more components onto integrated circuits. *Proc. IEEE* 86, 82–85.
35. Wright, T.P. (1936). Factors affecting the cost of airplanes. *J. Aeronaut. Sci.* 3, 122–128.
36. Thompson, P. (2012). The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing. *J. Econ. Perspect.* 26, 203–224.
37. Witajewski-Baltvilks, J., Verdolini, E., and Tavoni, M. (2015). Sending the learning curve. *Energy Econ* 52, S86–S99. <https://doi.org/10.1016/j.eneco.2015.09.007>.
38. Alberth, S. (2008). Forecasting technology costs via the experience curve—myth or magic? *Technol. Forecasting Soc. Change* 75, 952–983.
39. Lafond, F., Bailey, A.G., Bakker, J.D., Rebois, D., Zadorian, R., McSharry, P., and Farmer, J.D. (2018). How well do experience curves predict technological progress? a method for making distributional forecasts. *Technol. Forecasting Soc. Change* 128, 104–117.
40. IPCC (2022). *Mitigation of climate change. Contribution of working Group III to the sixth assessment report of the Intergovernmental Panel on Climate Change*. In *Climate Change* (Cambridge University Press).
41. Riahi, K., Kriegler, E., Johnson, N., Bertram, C., den Elzen, M., Eom, J., Schaeffer, M., Edmonds, J., Isaac, M., Krey, V., et al. (2015). Locked into Copenhagen pledges — implications of short-term emission targets for the cost and feasibility of long-term climate goals. *Technol. Forecasting Soc. Change* 90, 8–23. <https://doi.org/10.1016/j.techfore.2013.09.016>.
42. Krey, V., Guo, F., Kolp, P., Zhou, W., Schaeffer, R., Awasthy, A., Bertram, C., de Boer, H.-S., Fragkos, P., Fujimori, S., et al. (2019). Looking under the hood: a comparison of techno-economic assumptions across national and global integrated assessment models. *Energy* 172, 1254–1267. <https://doi.org/10.1016/j.energy.2018.12.131>.
43. Huppmann, D., Kriegler, E., Krey, V., Riahi, K., Rogelj, J., Calvin, K., Humpenoeder, F., Popp, A., Rose, S.K., Weyant, J., et al. (2019). °C *Scenario Explorer and Data Hosted by IIASA*. Technical report (IIASA).
44. Ferguson, C.D., Marburger, L.E., Doyne Farmer, J.D., and Makhijani, A. (2010). A US nuclear future? *Nature* 467, 391–393. <https://doi.org/10.1038/467391a>.
45. Wilson, C., Grubler, A., Bauer, N., Krey, V., and Riahi, K. (2013). Future capacity growth of energy technologies: are scenarios consistent with historical evidence? *Clim. Change* 118, 381–395. <https://doi.org/10.1007/s10584-012-0618-y>.
46. Anandarajah, G., McDowall, W., and Ekins, P. (2013). Decarbonising road transport with hydrogen and electricity: long term global technology learning scenarios. *Int. J. Hydr. Energy* 38, 3419–3432.
47. Heuberger, C.F., Rubin, E.S., Staffell, I., Shah, N., and Mac Dowell, N. (2017). Power capacity expansion planning considering endogenous technology cost learning. *Appl. Energy* 204, 831–845.
48. DeCarolis, J., Daly, H., Dodds, P., Keppo, I., Li, F., McDowall, W., Pye, S., Strachan, N., Trutnevte, E., Usher, W., et al. (2017). Formalizing best practice for energy system optimization modelling. *Appl. Energy* 194, 184–198.
49. Fricko, O., Havlik, P., Rogelj, J., Klimont, Z., Gusti, M., Johnson, N., Kolp, P., Strubegger, M., Valin, H., Amann, M., et al. (2017). The marker quantification of the shared socioeconomic pathway 2: a middle-of-the-road scenario for the 21st century. *Global Environ. Change* 42, 251–267. <https://doi.org/10.1016/j.gloenvcha.2016.06.004>.
50. Gritsevskiy, A., and Nakićenovi, N. (2000). Modeling uncertainty of induced technological change. *Energy Policy* 28, 907–921.
51. Trutnevte, E. (2016). Does cost optimization approximate the real-world energy transition? *Energy* 106, 182–193.
52. Davis, S.J., Lewis, N.S., Shaner, M., Aggarwal, S., Arent, D., Azevedo, I.L., Benson, S.M., Bradley, T., Brouwer, J., Chiang, Y.-m., et al. (2018). Net-zero emissions energy systems. *Science* 360, eaas9793. <https://doi.org/10.1126/science.aas9793>.
53. Brown, T., Schlachtberger, D., Kies, A., Schramm, S., and Greiner, M. (2018). Synergies of sector coupling and transmission reinforcement in a cost-optimised, highly renewable European energy system. *Energy* 160, 720–739. <https://doi.org/10.1016/j.energy.2018.06.222>.
54. Eyre, N. (2019). *Energy efficiency in the energy transition* (European Council for an Energy Efficient Economy Summer Study Proceedings), pp. 247–254.
55. McNerney, J., Doyne Farmer, J., and Trancik, J.E. (2011). Historical costs of coal-fired electricity and implications for the future. *Energy Policy* 39, 3042–3054.
56. Krey, V. (2014). Global energy-climate scenarios and models: a review. *WIREs Energy Environ* 3, 363–383. <https://doi.org/10.1002/wene.98>.
57. Edenhofer, O., Lessmann, K., Kемfert, C., Grubb, M., and Köhler, J. (2006). Induced technological change: exploring its implications for the economics of atmospheric stabilization: synthesis report from the innovation modeling comparison project. *Energy J*. SI2006, 19449089. <http://www.jstor.org/stable/23297057>.
58. Nordhaus, W.D. (2017). Revisiting the social cost of carbon. *Proc. Natl. Acad. Sci. USA* 114, 1518–1523.
59. Broome, J. (2008). The ethics of climate change. *Sci. Am.* 298, 96–102. <http://www.jstor.org/stable/26000646>.
60. Stern, N. (2007). *The Economics of Climate Change: the Stern Review* (Cambridge University Press).
61. Woo, J.R., and Magee, C.L. (2021). Relationship between technological improvement and innovation diffusion: an empirical test. *Technol. Anal. Strateg. Manag.* 1–16. <https://doi.org/10.1080/09537325.2021.1901875>.
62. Lowe, R.J., and Drummond, P. (2022). Solar, wind and logistic substitution in global energy supply to 2050 – barriers and implications. *Renew. Sustain. Energy Rev.* 153, 111720. <https://doi.org/10.1016/j.rser.2021.111720>.
63. Mercure, J.-F., Pollitt, H., Viñuales, J.E., Edwards, N.R., Holden, P.B., Chewprecha, U., Salas, P., Sognaes, I., Lam, A., and Knobloch, F. (2018). Macroeconomic impact of stranded fossil fuel assets. *Nat. Clim. Change* 8, 588–593. <https://doi.org/10.1038/s41558-018-0182-1>.
64. McNerney, J., Farmer, J.D., Redner, S., and Trancik, J.E. (2011). Role of design complexity in technology improvement. *Proc. Natl. Acad. Sci. USA* 108, 9008–9013.
65. Wilson, C., Grubler, A., Bento, N., Healey, S., De Stercke, S., and Zimm, C. (2020). Granular technologies to accelerate decarbonization. *Science* 368, 36–39. <https://doi.org/10.1126/science.aaz8060>.
66. Verdolini, E., Anadón, L.D., Baker, E., Bosetti, V., and Aleluia Reis, L.A. (2018). Future prospects for energy technologies: insights from expert elicitations. *Rev. Environ. Econ. Policy* 12, 133–153. <https://doi.org/10.1093/reep/rev028>.
67. Meng, J., Way, R., Verdolini, E., and Diaz Anadon, L.D. (2021). Comparing expert elicitation and model-based probabilistic technology cost forecasts for the energy transition. *Proc. Natl. Acad. Sci. USA* 118, e1917165118. <https://doi.org/10.1073/pnas.1917165118>.
68. Zimm, C., Goldemberg, J., Nakićenovi, N., and Busch, S. (2019). Is the renewables transformation a piece of cake or a pie in the sky? *Energy Strategy Rev* 26, 100401. <https://doi.org/10.1016/j.esr.2019.100401>.
69. Grubb, M., Wieners, C., and Yang, P. (2021). Modeling myths: on dice and dynamic realism in integrated assessment models of climate change mitigation. *WIREs Clim. Change* 12, e698. <https://doi.org/10.1002/wcc.698>.
70. Köberle, A.C., Vandyck, T., Guivarch, C., Macaluso, N., Bosetti, V., Gambhir, A., Tavoni,

- M., and Rogelj, J. (2021). The cost of mitigation revisited. *Nat. Clim. Change* *11*, 1035–1045. <https://doi.org/10.1038/s41558-021-01203-6>.
71. Luderer, G., Madeddu, S., Merfort, L., Ueckerdt, F., Pehl, M., Pietzcker, R., Rottoli, M., Schreyer, F., Bauer, N., Baumstark, L., et al. (2022). Impact of declining renewable energy costs on electrification in low-emission scenarios. *Nat. Energy* *7*, 32–42. <https://doi.org/10.1038/s41560-021-00937-z>.
 72. Way, R., Lafond, F., Lillo, F., Panchenko, V., and Farmer, J.D. (2019). Wright meets Markowitz: how standard portfolio theory changes when assets are technologies following experience curves. *J. Econ. Dyn. Control* *101*, 211–238. <https://doi.org/10.1016/j.jedc.2018.10.006>.
 73. Pindyck, R.S. (2019). The social cost of carbon revisited. *J. Environ. Econ. Manag.* *94*, 140–160. <https://doi.org/10.1016/j.jeem.2019.02.003>.
 74. Kikstra, J.S., Waidelich, P., Rising, J., Yumashev, D., Hope, C., and Brierley, C.M. (2021). The social cost of carbon dioxide under climate-economy feedbacks and temperature variability. *Environ. Res. Lett.* *16*, 094037. <https://doi.org/10.1088/1748-9326/ac1d0b>.
 75. Aldy, J.E., Kotchen, M.J., Stavins, R.N., and Stock, J.H. (2021). Keep climate policy focused on the social cost of carbon. *Science* *373*, 850–852. <https://doi.org/10.1126/science.abi7813>.
 76. IEA (2019). *Offshore wind Outlook 2019*. Technical report (International Energy Agency).
 77. De Stercke, S. (2014). Dynamics of Energy Systems: a Useful Perspective. IIASA interim report (IIASA). <http://pure.iiasa.ac.at/id/eprint/11254/>.
 78. Ziegler, M.S., and Trancik, J.E. (2021). Re-examining rates of lithium-ion battery technology improvement and cost decline. *Energy Environ. Sci.* *14*, 1635–1651. <https://doi.org/10.1039/D0EE02681F>.
 79. IEA (2019). *The Future of Hydrogen*. Technical report (International Energy Agency).
 80. Tong, D., Farnham, D.J., Duan, L., Zhang, Q., Lewis, N.S., Caldeira, K., and Davis, S.J. (2021). Geophysical constraints on the reliability of solar and wind power worldwide. *Nat. Commun.* *12*, 6146. <https://doi.org/10.1038/s41467-021-26355-z>.
 81. IEA (2021). *World Energy Outlook 2021*. Technical report (International Energy Agency).
 82. BP (2021). *BP Statistical Review of World Energy*. Technical report (BP plc).
 83. IRENA (2021). *Renewable Power Generation Costs in 2020* (International Renewable Energy Agency). Technical report.
 84. Lazard (2021). *Lazard's Levelized Cost of Energy Analysis Versions 2–15*. Technical report (Lazard).
 85. EIA (2021). *Annual Energy Outlook Reports 1979–2021*. Technical report (Energy Information Administration).

Joule, Volume 6

Supplemental information

**Empirically grounded technology
forecasts and the energy transition**

Rupert Way, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer

Supplemental Experimental Procedures

Contents

List of symbols and abbreviations	3
1 Model description	5
1.1 General approach	5
1.2 Energy system description	6
1.3 Model components	7
1.4 Final energy components included in the model	12
1.4.1 Oil, coal, gas and electricity	12
1.4.2 Energy requirements of the energy sector itself	13
1.5 Final energy components excluded from the model	13
1.5.1 Heat (used as an energy carrier)	13
1.5.2 Traditional Biomass	13
1.5.3 Petrochemical feedstock	13
1.5.4 Bioenergy, solar thermal energy, marine energy and geothermal energy	14
1.6 Other energy system components	14
1.6.1 CCS	14
1.6.2 Electricity technologies and networks	15
1.6.3 P2X fuels	15
1.6.4 Fuels infrastructure	16
1.6.5 Electrification of transport	17
1.6.6 Other demand-side capital costs	17
1.7 Units and justification for the use of LCOE	17
1.8 End-use conversion efficiencies	19
1.9 Estimating total system costs	21
2 Scenario construction	26
2.1 Growth rate model	26
2.2 Energy services	28
2.3 Indices, slack variables and meeting the useful energy constraint	29
2.4 Energy carrier quantities for sectors	30
2.5 Electricity generation technologies	32
2.6 Energy sector fossil fuels	32
2.7 Direct-use fossil fuels	33
3 Energy storage and flexibility requirements	34
3.1 Short-term batteries	37
3.2 Multi-day storage	38
3.3 P2X fuels for long term energy storage and hard-to-electrify applications	38
3.3.1 P2X fuel storage	39
3.3.2 P2X fuel transmission and distribution	40
3.3.3 P2X fuel quantity	40
3.4 Verifying electricity system reliability	41
3.5 Electricity adjustment for hydrogen	43
3.6 Electrolyzers	43
3.7 Electricity networks	44
3.7.1 Grid flexibility	44

3.7.2	Grid size	45
3.8	Technology lifespan and annual additions	49
4	Scenarios	51
4.1	No Transition	53
4.2	Fast Transition	56
4.3	Slow Transition	59
4.4	Slow Nuclear Transition	62
4.5	Historical Mix	65
4.6	Comparison with AR6 scenarios	68
5	Technology cost models	72
5.1	AR(1) process	72
5.2	Wright's law	73
5.3	Total costs of AR(1) and Wright's law technologies	74
5.3.1	AR(1) model costs	74
5.3.2	Wright's law model costs	74
5.4	Total system costs	75
5.5	Net present cost of transition	76
6	Data, calibration and technology forecasts	77
6.1	Data sources for Figure 1	77
6.2	Oil	80
6.3	Coal	81
6.4	Gas	82
6.5	Coal electricity	83
6.6	Gas electricity	85
6.7	Nuclear electricity	87
6.8	Hydropower electricity	91
6.9	Biopower electricity	93
6.10	Wind electricity	95
6.11	Solar PV electricity	97
6.12	Batteries	100
6.13	Hydrogen and electrolyzers	103
6.14	Electricity networks	108
6.15	Technology data summary	109
7	Results	111
7.1	Main case results	111
7.2	Sensitivity to time horizon T	113
7.3	Sensitivity to system growth rate	114
7.4	Sensitivity to cost assumptions	119
7.4.1	Less pessimistic solar costs	119
7.4.2	Lower oil and gas prices	121
7.4.3	Lower nuclear costs	122
7.4.4	Less uncertain electrolyzer costs	124
7.5	Moore's law results	126
8	Further Discussions	131
8.1	The heterogeneity and persistence of technological change	131
8.2	Bias-variance trade-off	132
8.3	Wright's law caveats	133
8.3.1	Discussion on questions of causality	133
8.3.2	Further caveats	135
8.4	Comparison to existing energy system modelling approaches	136
8.5	The use of floor costs in endogenous technological learning models	138
8.6	Regional differences in competing technologies	139

8.7	Is the speed of the Fast Transition achievable?	142
8.7.1	Are the growth rates of electrolyzers in the Fast Transition plausible?	143
8.8	Additional benefits from the Fast Transition	146
8.8.1	Emission reductions and reduced climate risks	146
8.8.2	Co-benefits of a Fast Transition	148
8.9	Additional costs from the Fast Transition	149
8.9.1	Stranded Assets	149

List of symbols and abbreviations

Abbreviations

AEL	Alkaline electrolysis
BAU	Business-as-usual
bbbl	Barrel of crude oil
CCGT	Combined-cycle gas turbine
CCS	Carbon capture and storage
CO ₂	Carbon dioxide
COP	Coefficient of performance
CRF	Capital recovery factor
EOR	Enhanced oil recovery
EV	Electric vehicle
GHG	Greenhouse gas
Gt	Gigatonne
ICEV	Internal combustion engine vehicle
LCOE	Levelized cost of electricity
LHV	Lower heating value
Li-ion	Lithium-ion
LDES	Long-duration energy storage
MMBtu	Million British thermal units
NPC	Net present cost
PEM	Polymer-electrolyte membrane
PPA	Power purchase agreement
PtG	Power-to-gas
PtL	Power-to-liquids
P2X or PtX	Power-to-X
PV	Photovoltaic
SCC	Social cost of carbon
SMR	Steam methane reforming
SOEC	Solid oxide electrolyzer cell
SOFC	Solid oxide fuel cell
T&D	Transmission and distribution
TN or tn	Trillion
VRE	Variable renewable energy
VRF	Vanadium redox flow
yr	Year

Variables, constants and parameters

t_0	Base year
T	Number of periods in model
$t_0 + T$	Time horizon
t	Period after base year, in years
i	Technology index
j	Energy carrier index
k	Sector index
\mathcal{I}_j	Set of indices of technologies used as input to carrier j
\mathcal{J}_k	Set of indices of carriers used in sector k
i_{slack}	Index of slack electricity generation technology

j_{slack}^k	Index of slack carrier in sector k
$EF_t^{j,k}$	Final energy provided by carrier j to sector k in year t
$EF_t^{total,k}$	Total final energy provided by all carriers to sector k in year t
$EF_t^{j, total}$	Total final energy provided by carrier j to all sectors in year t
$EU_t^{j,k}$	Useful energy provided by carrier j in sector k in year t
$EU_t^{total,k}$	Total useful energy provided by sector k in year t
$EU_t^{j, total}$	Total useful energy provided by carrier j in year t
$Elec_t$	Total quantity of electricity generated in year t
$P2X\ fuels_t$	Total quantity of P2X fuels produced in year t
η_j^k	Conversion efficiency of final to useful energy of carrier j when used in sector k
$g_t^{j,k}$	Growth rate of carrier j in sector k in year t
$\xi^{j,k}$	Maximum percentage of useful energy from sector k that carrier j may provide at any time
q_t^i	Production of technology i in year t
g_t^i	Growth rate of technology i in year t
ψ^i	Maximum percentage of total electricity that may be generated by technology i at any time
β_{LT}	Fraction of electricity produced every day that multi-day storage must be able to store
β_{ST}	Fraction of electricity produced every day that short-term batteries must be able to store
g_{LT}^{grid}	Initial constant annual growth rate of multi-day storage used in power grid
g_{ST}^{grid}	Initial constant annual growth rate of short-term batteries used in power grid
$g_{ST}^{transport}$	Initial constant annual growth rate of short-term batteries used in transport sector

1 Model description

1.1 General approach

The purpose of this model is to *i*) generate probabilistic forecasts of the relative net costs associated with various energy system scenarios, and *ii*) determine which scenario is worth pursuing based on these forecast costs alone. To do this we use observed historical technology data and empirically validated cost forecasting methods to generate probability distributions of the future costs of single technologies, then combine these to estimate total energy system costs in various scenarios. Because it is not possible to forecast the costs of all energy system components using these forecasting methods (since sufficient data for model calibration is not always available), we apply some “worst-case” assumptions that allow us to bound the *relative* expected net present cost (NPC) of different scenarios. This reveals the preferred scenario in terms of expected costs. We also compute the probability that this scenario will be cheaper than its competitors, i.e. the probability that pursuing this scenario will turn out to be a winning bet.

Our perspective is that since the future is intrinsically uncertain, decision making is always a gamble, and so the most important determinants of the future development of the energy system – arguably technology costs – should be characterized probabilistically, using empirically validated models calibrated with observed data. We thus reformulate the energy transition question in terms of probabilistic technology cost forecasts. The evidence for induced innovation in energy technologies and systems is strong¹, so we use an empirically validated model embodying this dynamic – a stochastic experience curve model developed and tested by Lafond et al.² – to forecast key technology costs. In this framework, different energy system scenarios are viewed as different bets on future technological progress, conditional upon observed past progress trends. Our model answers the question: given the consistent cost reductions in key clean energy technologies observed over the last several decades, what is the rough shape (in terms of time-scales and energy mix) of the energy system scenario that is likely to be cheapest? The model only considers costs directly associated with supplying energy services to the global economy, and does not consider any benefits (or costs) related to health, ecosystems, welfare, equality, migration, trade, systemic risk, etc.

We use the current energy system as a benchmark scenario, and contrast its forecast cost against the forecast costs of two alternative scenarios that implement transitions to a clean energy system. The energy system model we construct is designed to only represent explicitly the very largest energy flows and most important technology substitutions, and neglect everything that, from our current perspective, appears negligible. Other essential components of the energy system (e.g. technologies that facilitate the transport and distribution of energy), are represented either explicitly, when possible, or implicitly by some simple worst-case-assumption bounds. This strategy allows us to determine relative scenario costs.

We model the core components of the global energy system each year from 2021 to 2070 using 14 supply-side energy technologies: 3 direct-use fossil fuels, 7 electricity generators, 3 energy storage technologies (two types of battery, plus electrolyzers), and the power grid. The model is global and does not differentiate between different geographic regions. The high level of aggregation in all modelling dimensions (regions, technology types, technology vintages etc.) is carefully chosen to ensure that the modelled system is represented in a manner appropriate for the probabilistic cost forecasting methods used. Lower levels of aggregation are possible, but would obscure the clarity of the presentation here.

The benchmark scenario is called *No Transition*, and the two others are called *Fast Transition* and *Slow Transition*. These are shown in detail in S.I. 4, but we also refer to them repeatedly before that in order to explain how scenarios are constructed and how total system costs are estimated. In brief, *No Transition* is a scenario in which unabated fossil fuel consumption remains at its current level, so that the economy relies primarily on fossil fuels far in to the future; *Fast Transition* is a scenario in which the current high growth rates of clean energy technologies persist in to the future, so that fossil fuels are displaced within about two decades; and *Slow Transition* is a scenario between these two, with clean energy growth rates approximately halving, then persisting in to the future, so that fossil fuel displacement is slow until 2040 then gathers pace. (Two further scenarios, called *Slow Nuclear Transition* and *Historical Mix*, are also presented in S.I. 4, but are not featured prominently in the paper because they turned out to be largely irrelevant to the main results.) We do not claim that any of these is the cheapest scenario possible, we simply use them to compare the effects of different types and speeds of transition on relative scenario costs.

There are two steps to the overall modelling procedure. First, the model energy system scenarios are constructed, subject to the constraints that they must all provide identical levels of energy services to all end-use sectors of the economy, and satisfy all expected engineering constraints. There are no market mechanisms involved in the construction process: scenarios are created by specifying exogenously varying technology

growth rates. Second, empirically validated probabilistic technology cost forecasting methods are calibrated using extensive historical data, and used to forecast individual technology cost distributions in each scenario. These are summed and discounted to give probability distributions of the NPC of each scenario, from which the expected NPC of transition for each scenario, relative to the No Transition scenario, is computed.

This approach is similar to that employed by Gritsevskiy and Nakicenovic³. The most significant differences are our carefully validated and calibrated cost forecasting models, and our decision to develop a medium-complexity model (as opposed to their high complexity model with many more technologies and processes). This is in order to avoid many of the pitfalls commonly observed in both low and high complexity models^{4,5,6,7}, including: too much model granularity to understand the main drivers of results; too little model granularity to capture the diversity of technological progress; excessively long run-times, which prevent detailed exploration of parameter space and parameter sensitivity; difficulty of calibration and keeping the model up to date with the latest data; difficulty of model validation (backtesting); and the inability to represent feedback loops, resulting in systematic biases and large forecast errors for fast-progressing technologies.

As discussed further in S.I. 8.4, our simple two-step approach contrasts with most other energy system models, which usually involve modelling energy technologies and services, natural resources, infrastructure and geographical regions at a much higher resolution. As we emphasise in the main text, we have designed our model to be as parsimonious as possible while ensuring all assumptions and dynamics are both technically feasible and tied as closely as possible to observed time-series data. The simplicity and transparency of the model and its results make it easy to be understood, replicated, and modified, with only a relatively small number of assumptions.

We only include major energy flows in the model, and omit any minor energy flows that are not large enough to influence the results, because this allows us to focus on a few critical technology substitutions that can have a major impact on the energy system. We only explicitly model key features of the energy transition that differ dramatically between scenarios, and for everything else we use simple worst-case cost assumptions that represent the maximum estimated difficulty in transitioning away from the current system, though in all cases these are highly conservative and strongly favour the No Transition scenario (in terms of estimated relative costs). Thus our final results are in fact *upper bounds* on the relative net present costs of each scenario, relative to the No Transition baseline.

1.2 Energy system description

We describe the global energy system as follows (see GEA⁸ for comprehensive descriptions):

- *Primary energy* is the energy embodied in resources as they exist in nature. This metric is useful for describing physical substances that store energy in chemical bonds, such as fossil fuels and biomass-derived fuels. It is not very useful for describing other sources of energy though, such as nuclear, geothermal, solar or wind power, since the extent of the original energy source in these cases is not easily defined.
- *Energy carriers* are substances or phenomena in to which primary energy resources are converted for consumption by energy users in the economy. Electricity, gasoline, coal, and power-to-X (P2X) fuels are examples of energy carriers. Primary energy resources are generally converted to energy carriers to facilitate transport, storage, standardisation, and market transactions of energy.
- *Final energy* is the energy embodied in energy carriers as they are delivered to energy users in the economy. In this model, we define final energy to be the energy delivered to end-use sectors of the economy, *and the energy delivered to the energy sector for self-consumption*.*
- *End-use applications* (or *end-use technologies*) are the last step in the energy delivery chain, where energy carriers are used by a consumer to perform some energy service.
- *Energy services* are services provided to consumers. The main energy services are heating/cooling, mobility, consumer goods and food. Estimating the minimum quantity of energy required to provide an energy service is not at all straightforward, and varies over time as new technologies are developed

*Note the distinction between this and a perhaps more common definition, in which final energy refers only to energy carriers delivered to end-use sectors of the economy, but not to the energy sector itself. The definition we use has the benefit that it allows energy consumption of the energy sector to be described and modelled in the same way as other sectors, which is useful when considering the large scale substitution processes involved in the energy transition.

that are able to use energy more efficiently. In many cases there is a large efficiency gap between this minimum energy requirement and the energy consumed in practise.

- *Useful energy* is the portion of final energy that goes towards doing the useful work that energy services provide, while the rest is lost to the environment. It is the energy required by energy technologies, in various forms such as kinetic energy, heat and light, to perform energy services. End-use application technology choices have a large impact on both useful and final energy consumption because their conversion efficiencies vary widely. For example, consider providing an equal level of transport service via an electric vehicle (EV) or an internal combustion engine vehicle (ICEV). The amount of useful energy consumed in both cases is the same. But while the EV is around 80% efficient at converting final energy (electricity) to useful energy (kinetic energy), the ICEV is only around 25% efficient (75% of the chemical energy in petrol is wasted as heat), so requires around three times as much final energy. Due to wasted heat, solid, liquid and gaseous fuels usually have very large discrepancies between the final energy they contain and the useful energy they provide.
- *Energy technologies*, as commonly understood, are any technologies involved in obtaining, converting, storing, distributing or using energy. The precise meaning of the phrase is not universally agreed upon, and it is applied flexibly in different contexts. In this model we use a very specific definition. Energy technologies *in this model* (also referred to as simply *technologies* henceforth) are defined to be precisely those components of the model whose costs are specified and included in the total system cost calculation. We use this terminology because most model components' costs are obtained via a *technology cost forecasting model*.*

While we note that the concept of primary energy is frequently used by energy agencies, companies, and modellers, it unfortunately obscures important information about conversion efficiencies and the energy wasted in fuel consumption. In this paper we focus on final and useful energy instead.

1.3 Model components

The goal of the work here is to use data-driven technology forecasts to estimate total system costs associated with various energy scenarios, so a critical bottleneck is the availability of historical cost (or price) data for energy technologies, which is used to calibrate the forecast models. The technology cost forecasting methods we use apply to single technologies (see S.I. 5), so to investigate the impact of these forecasts on total system costs, we need some method for combining technologies together in a model system that is as close to the real energy system as possible, and then exploring the evolution of this system through time. We construct a model energy system that is able to draw on as much historical data as possible, and hence capture long-term trends in technological progress as accurately as possible. At the same time, however, the model must be sufficiently technically realistic, and represent as much of the modelled system as possible.

For crude oil, coal, natural gas, and coal electricity, consistent annual price data exists for around a century. For nuclear, solar and wind electricity, costs are available for several decades, generally beginning in years from the 1970s to the 90s. Cost data for other electricity generation technologies, batteries and electrolyzers are typically only available for a few decades. For most other energy technologies only very limited or irregular data exists, making it hard to construct satisfactory cost forecasts.

The key data around which the model is constructed are:

- Price of crude oil (\$/bbl), natural gas (\$/MMBtu) and coal (\$/short ton) (long time series).
- Levelized cost of electricity (LCOE) (\$/MWh) generated from coal, gas, nuclear, hydropower, biopower, wind, solar photovoltaics (PV) (medium or short time series).
- Capital cost of: lithium-ion (Li-ion) batteries (\$/kWh), Vanadium redox flow batteries (\$/kWh) and electrolyzers (\$/kW) (medium or short time series).

Historical cost (or price) and production data for all these energy technologies exist, so their future costs may be forecast based on relatively good historical data. This is not the case for other technologies. Our energy

*While this definition coincides with the familiar understanding of energy technologies in most cases, there are a few idiosyncrasies. In this model, direct-use oil, coal and gas are all considered to be "energy technologies" in their own right, but P2X fuels are not. This is because the costs of direct-use oil, coal and gas are included explicitly in the total system cost calculation, but P2X fuel cost is only accounted for *indirectly*, as the sum of its constituents' costs: electricity generation and installed electrolyzer capacity. Thus in our terminology, P2X fuels are considered an energy carrier, but not an energy technology.

system model is therefore designed to maximise the fraction of the real-world energy system that can be accurately represented by some configuration of just these 13 technologies (plus power grid costs, which are dealt with separately). Data and variable choices are discussed further in the rest of this section as well as S.I. 5 and S.I. 6. The components of our model energy system are:

- 3 end-use sectors of the economy: transport, industry and buildings.
- 1 intermediate sector: energy.
- 5 energy carriers: oil, coal, gas, electricity, P2X fuels.
- 10 energy inputs:
 - 3 direct-use primary energy resources: crude oil, coal and gas.
 - 7 electricity generation technologies: coal, gas, nuclear, hydropower, biopower, wind and solar.
- 3 storage and conversion technologies: daily-cycling batteries, multi-day storage, electrolyzers. Daily batteries are used in both the intermediate and transport sectors.
- 1 energy distribution technology: electricity networks (i.e. the power grid).

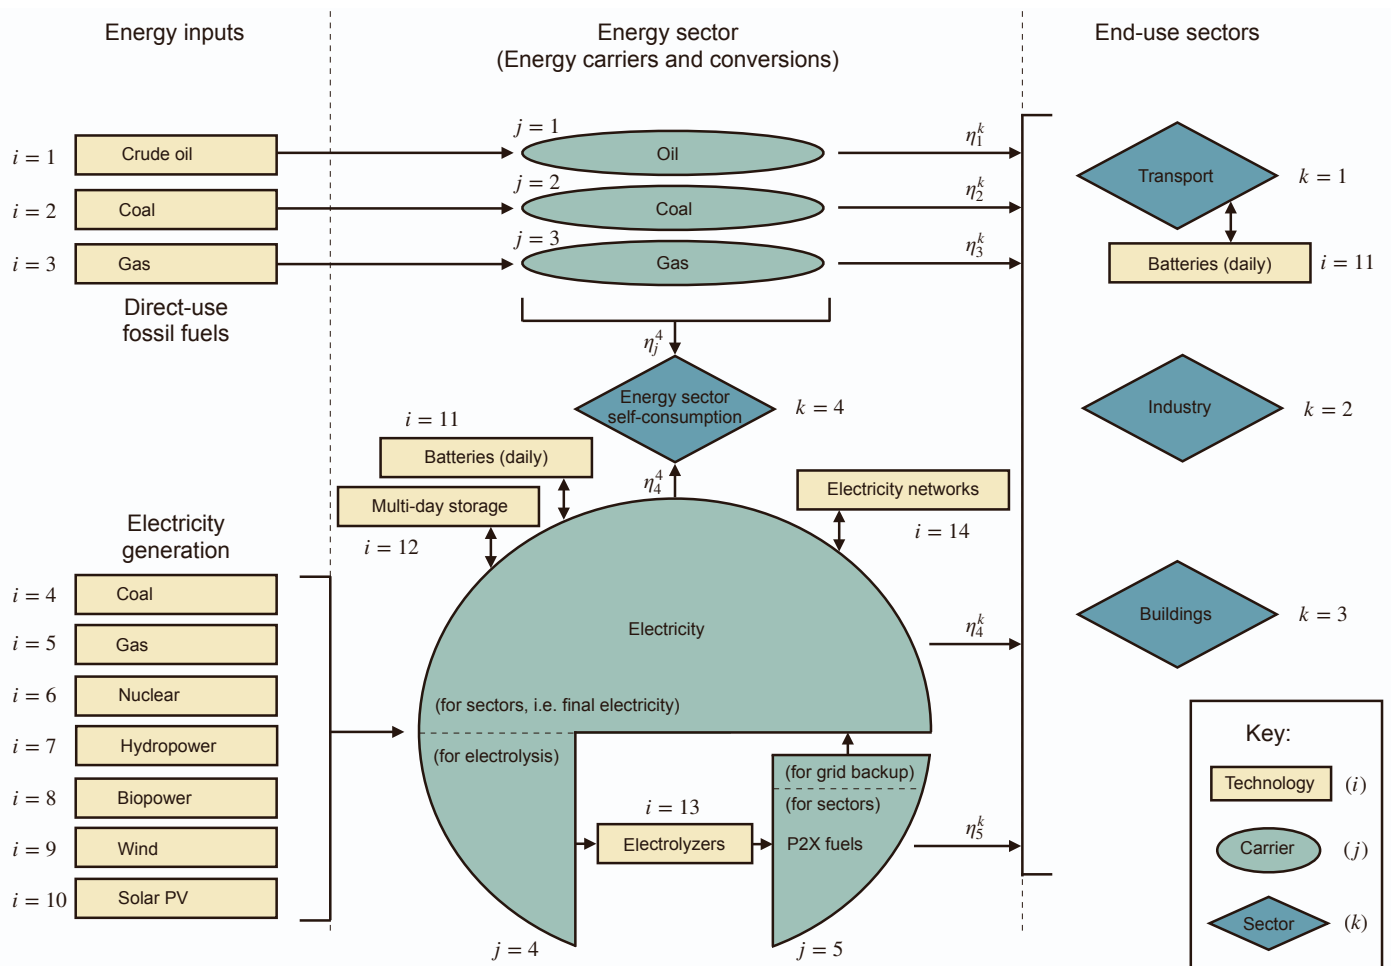


Figure S1: Energy system model components. Historical costs and production are known for all *Technologies*, so their future costs may be modelled. Scenarios are constructed by specifying suitable quantities of all components to meet sector requirements. The system cost is then estimated by summing up forecast costs of all technologies in a given scenario. Indices i, j, k specify technologies, energy carriers, and sectors, respectively, and are described in S.I. 2.3. The conversion factors η_j^k represent the efficiency with which final energy is converted to useful energy, for each carrier/sector pair, as described in S.I. 1.8.

The model system is shown in Figure S1. It covers as much of the present and potential future energy system as possible using as few components (with sufficient historical data) as possible. End-use sectors and the energy sector provide energy services to consumers, be they individuals, firms, or institutions. To do this they

receive energy carriers from the energy sector, and use the energy contained therein to perform useful work. In this model we refer to the energy content of energy carriers as “final energy” for all sectors, including the energy sector.* Thus all sectors receive final energy and convert it to useful energy to provide energy services. The energy sector takes energy inputs and converts them to energy carriers; in doing so it self-consumes some energy carriers. Energy technologies provide all supply, storage, conversion and distribution capabilities required to produce and manage energy carriers throughout the entire system. At present the energy system has almost no P2X fuels or batteries, and very large quantities of fossil fuels, but this would be reversed in the clean energy scenarios considered here.

An important feature of our system model is that we consider all electricity generators to be energy sources in their own right. Many energy models include fossil fuel electricity generation within the energy sector, and model the full production pathway from raw materials to electricity. In contrast, when electricity is used, we do not consider processes prior to the electricity generation stage, in particular the production, transport, and conversion of fossil fuels. This is because the costs of all these processes are ultimately captured in LCOE values, so by focusing on just the electricity generation stage, our method allows us to use historical LCOE data series to estimate future system costs, while also avoiding the need to specify large numbers of additional parameters and detailed conversion processes.

Our system model is designed in such a way that all four sectors of the economy are treated identically in the scenario modelling process. Each sector receives final energy, in the form of energy carriers, and uses it to process various other (unspecified) inputs, and perform useful tasks. Energy carriers can be substituted for one another, subject to basic engineering constraints, to model decarbonization pathways. Note that end-use sectors receive all five energy carriers (P2X fuels are required because they can substitute for fossil fuels in some end-use applications), but the energy sector only receives oil, coal, gas and electricity. This is because the energy sector has no need for P2X fuels – the fossil fuel energy carriers are used to process other fossil fuels for end-use consumption, and so any transition pathway that removes fossil fuels from end-use sectors automatically removes the need for fossil fuels (or P2X fuel substitutes) in the energy sector too.

As well as receiving final energy used to do useful work, the energy sector also receives “energy inputs” – crude oil, coal, gas, and seven types of electricity generation. These are simply considered as process inputs, to which energy sector processes are applied to produce commercial-grade energy carriers, which are then shipped throughout the economy. In this framework, as well as taking crude oil, coal, and gas, and producing the corresponding energy carriers, the energy sector also takes electricity generation sources and produces both final electricity (accounting for all line losses) and P2X fuels (P2X fuel production requires no additional final energy inputs). As a result of these design choices, our system model accounts for all significant energy transformation losses, while also providing a concise description of energy carriers and sectors, which can be simply and systematically modelled.

We focus primarily on supply-side energy technologies in this work. There are many possibilities for innovation, efficiency increases and improvements in demand-side technologies that could further lower the costs of transition scenarios considered here^{9,10,11}. We exclude such considerations on the grounds of simplicity and communicability, since they are not required to demonstrate our main results.

Table S1 shows which electricity generation technologies are included in the model and which are omitted, and the quantities of electricity they generated in 2020, in exajoules (EJ).

*We use this definition primarily because it permits the most concise nomenclature when defining the model equations in S.I. 2.4. Whether one refers to energy carriers self-consumed by the energy sector as final energy or not is simply a labelling issue, and has no bearing on the results.

Technology	Electricity generation, EJ	
	Included	Excluded
Oil		2.7
Coal	35.7	
Gas	22.9	
Nuclear	10.0	
Hydro	15.2	
Biopower	2.55	
Wind	5.75	
Geothermal		0.34
Solar PV	3.0	
CSP		0.047
Marine		0.004
Totals	95.1	3.091
	96.9%	3.1%

Table S1: Electricity generation in 2020 (or 2019) by technology, as listed in the International Energy Agency’s (IEA’s) 2021 World Energy Outlook (WEO)¹², showing whether the technology is included or excluded from our model. (Values from 2019 are used when the 2020 values show a substantial single-year decrease due to the COVID-19 pandemic.) (Column sums may differ from totals due to rounding.)

Table S2 gives a detailed breakdown of which energy flows are included in the model, and which are omitted, in exajoules. For convenience, Table S3 provides a copy of Table S2 in units of Mtoe, allowing for direct comparison with the original data source (Table A.3 of the IEA’s WEO 2019¹³). Individual components in the table are discussed in the subsections below.

The energy system components included in the model currently account for around 97% of electricity generation and 95% of CO₂ emissions from the power sector; and 83% of all final energy supplied to end-use sectors and 82% of the related CO₂ emissions. However, several of the minor energy carriers not included in the model are renewable technology energy carriers, namely: bioenergy, biofuels, traditional biomass, and “other renewables”. Note that these are energy *carriers*, so biopower is not included in this group (and is indeed included in the model). In the long run, these carriers do not produce *net* CO₂ emissions (depending on how they are sourced), so are not in the same qualitative category as fossil fuel technologies, and can be omitted without affecting model results or the main conclusions regarding decarbonisation. Petrochemical feedstock can similarly be omitted from the model, as it does not produce combustion emissions (see S.I. 1.5.3). Hence, the most appropriate way to view the model is that it covers 90% of all the final energy components that produce significant net CO₂ emissions, as shown in Row block C of Table S2. In summary, *the model covers 90% of all current final energy, excluding the renewable energy carriers listed above and petrochemical feedstock.*

Energy carrier	End-use sectors						Energy sector, EJ		"Other" uses, EJ		Totals by carrier, EJ	
	Transport, EJ		Industry, EJ		Buildings, EJ		Incl.	Excl.	Incl.	Excl.	Incl.	Excl.
	Incl.	Excl.	Incl.	Excl.	Incl.	Excl.						
<i>Included in model:</i>												
Direct-use oil	110.1		12.5		13.8							136.4
Direct-use coal			33.3		5.2							38.5
Direct-use gas			27.0		29.3							56.3
Electricity	1.3		33.6		42.3		15.5					92.8
Unspecified fossil fuels							47.5					47.5
<i>Excluded due to small size:</i>												
Heat (used as energy carrier)				5.99		6.28						12.3
Bioenergy (excl. Trad. Biom.)*				8.92		4.9						13.8
Biofuels*		3.7										3.7
Other fuels		4.7										4.7
Other renewables*				0.04		2.0						2.0
<i>Excluded as non-comparable:</i>												
Traditional biomass*						26.0						26.0
<i>Excluded as unknown:</i>												
Unspecified										22.7		22.7
<i>Excluded as non-energy:</i>												
Petrochem. feedstock										(23.0)		(23.0)
A. Entire system final energy (excl. petrochem. feedstock)												
Column totals, EJ	111.4	8.4	106.4	15.0	90.7	39.2	63.0	0.0	0.0	22.7	371.5	85.2
Sector totals, EJ	120		121		130		63		23		457	
Model coverage	93%		88%		70%		100%		0%		81%	
B. End-use sector final energy coverage												
End-use modelled, EJ			309									
End-use total, EJ			371									
Model coverage			83%									
C. System final energy excl. renewable energy carriers (*) (and petrochem. feedstock)												
Column totals, EJ	111.4	4.7	106.4	5.99	90.7	6.28	63.0	0.0	0.0	22.7	371.5	39.7
Sector totals, EJ	116		112		97		63		23		411	
Model coverage	96%		95%		94%		100%		0%		90%	

* Renewable technology energy carrier

Table S2: Energy uses in 2018 by sector and carrier, as listed in Table A.3 of the IEA's WEO 2019¹³. Each column for a given sector is split into two sub-columns, one showing the components that are included in our model, the other showing those that are excluded. Row block A shows the coverage of all modelled final energy components, relative to the entire energy system (excluding petrochemical feedstock). Row block B shows the fraction of end-use sector final energy covered by the model. Row block C shows the coverage of all modelled components, relative to the energy system excluding all minor energy carriers that are renewable technologies already, shown by asterisks (and again excluding petrochemical feedstock). Petrochemical feedstock values are in brackets as a reminder that they are excluded from all energy carrier totals. Row and column sums differ from totals due to rounding. Note that this table only shows final energy flows; storage, transport, distribution and conversion technologies are not listed here. See Table S3 for values in units of Mtoe.

Energy carrier	End-use sectors						Energy sector, Mtoe		“Other” uses, Mtoe		Totals by carrier, Mtoe	
	Transp., Mtoe Incl.	Excl.	Ind., Mtoe Incl.	Excl.	Build., Mtoe Incl.	Excl.	Incl.	Excl.	Incl.	Excl.	Incl.	Excl.
<i>Included in model:</i>												
Direct-use oil	2629		298		330						3257	
Direct-use coal			795		125						920	
Direct-use gas			646		700						1346	
Electricity	32		803		1011		371				2217	
Unspecified fossil fuels							1135				1135	
<i>Excluded due to small size:</i>												
Heat (used as energy carrier)				143		150						293
Bioenergy (excl. Trad. Biom.)*				213		117						330
Biofuels*		89										89
Other fuels		113										113
Other renewables*				1		48						49
<i>Excluded as non-comparable:</i>												
Traditional biomass*						620						620
<i>Excluded as unknown:</i>												
Unspecified										543		543
<i>Excluded as non-energy:</i>												
Petrochem. feedstock										(549)		(549)
A. Entire system final energy (excl. petrochem. feedstock)												
Column totals, EJ	2661	202	2542	357	2166	935	1506	0.0	0.0	543	8875	2037
Sector totals, EJ	2863		2898		3101		1506		543		10912	
Model coverage	93%		88%		70%		100%		0%		81%	
B. End-use sector final energy coverage												
End-use modelled, EJ			7369									
End-use total, EJ			8862									
Model coverage			83%									
C. System final energy excl. renewable energy carriers (*) (and petrochem. feedstock)												
Column totals, EJ	2661	113	2542	143	2166	150	1506	0.0	0.0	543	8875	949
Sector totals, EJ	2774		2685		2316		1506		543		9824	
Model coverage	96%		95%		94%		100%		0%		90%	

* Renewable technology energy carrier

Table S3: Table S2 in units of Mtoe, for convenience in comparing with the original data source. See Table S2 for details.

1.4 Final energy components included in the model

1.4.1 Oil, coal, gas and electricity

In line with our approach of only considering the largest energy flows, the only energy carriers in the current energy system that we model are direct-use oil, coal and gas, plus electricity. Many other energy carriers are used, but we exclude these from the model because of their relatively small impact on the system. Electricity generation technologies included in the model are: coal, gas, nuclear, hydropower, biopower, wind and solar PV. Note that direct-use coal and gas ($i = 2, 3$ in Figure S1) are treated as different technologies from coal and gas electricity generation ($i = 4, 5$ in Figure S1). This distinction is required because their costs are accounted for using different methods.

Electricity costs are quantified in terms of LCOEs, and observed LCOE data is used for model calibration. The following factors are included in LCOE calculations for electricity generation:

- capacity factors
- capital costs in terms of materials, construction, installation, land, and decommissioning

- operating costs including fuel costs, fixed and variable operations and maintenance costs
- financing costs for the asset's lifetime

There is some variation in the precise detail of what is included in LCOE calculations found in the literature, with some also including taxes and subsidies, grid connection, or transmission costs, but the differences are usually small.

1.4.2 Energy requirements of the energy sector itself

The energy sector covers “the use of energy by transformation industries and the energy losses in converting primary energy into a form that can be used in the final consuming sectors. It includes losses by gas works, petroleum refineries, blast furnaces, coke ovens, coal and gas transformation and liquefaction. It also includes energy used in coal mines, in oil and gas extraction and in electricity and heat production.”*

This sector currently consumes around 63 EJ of energy per year. Of this, 47.5 EJ/yr, or about 10.4% of all final energy, is direct-use fossil fuel consumption, required for extracting and processing other fossil fuels for use in end-use sectors and the energy sector itself (e.g. converting crude oil to gasoline). Because end-use fossil fuel consumption varies a lot across the scenarios we consider, the future evolution of the energy sector – and the amount of energy it uses – will also differ between scenarios. We model this by assuming that energy sector fossil fuel energy carrier consumption scales with the quantity of fossil fuels used in end-use sectors in each scenario (both as direct-use fuels and as coal- and gas-fired electricity). In the Fast Transition scenario, for example, since most fossil fuels are displaced from end-use sectors by 2040, this means that energy sector fossil fuel consumption undergoes a corresponding decline. In the No Transition scenario though, end-use fossil fuel consumption remains high, and therefore so does energy sector fossil fuel consumption.

Electricity consumption of the energy sector is currently 15.5 EJ/yr, and around half of this is energy lost as heat from all the power lines worldwide[†]. In contrast to fossil fuel consumption, energy sector electricity consumption will continue to grow in all scenarios – we include this in the model.

1.5 Final energy components excluded from the model

1.5.1 Heat (used as an energy carrier)

Here heat refers to energy traded commercially by transporting physical substances (usually gases or liquids) at specific, desired temperatures (either hot or cold). Heat is produced from fuel combustion, nuclear reactors, geothermal sources, sunlight and various industrial processes. It is often a by-product of other useful services, such as power generation and industrial processes. For these reasons it cannot be straightforwardly included here in the same modelling framework as the other system components. In any case, heat (used as an energy carrier) is only 2.7% of final energy in the current energy system and so we omit it.

1.5.2 Traditional Biomass

Traditional biomass refers to fuelwood, charcoal, animal dung, and agricultural residues used for cooking and heating. Globally it remains a substantial end-use energy flow, currently accounting for 20% of final energy in the buildings sector (26 EJ of 130 EJ), i.e. 5.7% of total final energy (or around 5.8% of useful energy). Figure 1(b) makes it clear that traditional biomass has grown only slightly over the last 140 years and has declined over the last 20 years. Our model assumes that total useful energy supply will continue to rise at 2% per year, so that energy services double by 2055. If traditional biomass usage remains constant, it will decline to become only 2.9% of all energy services by then. Replacing traditional biomass use by modern energy services will depend on global societal changes that are beyond the scope of our modelling here, but in any case, will only have a very minor impact on future system costs, and so we omit it.

1.5.3 Petrochemical feedstock

Petrochemical feedstock refers to the 23 EJ of fossil fuels currently used as raw material each year to produce plastics and other industrial chemicals and products (e.g. lubricants). We exclude this from our model on the basis that the fossil fuels used for petrochemical feedstock are not acting as energy carriers in the usual

*This is an extract from the Definitions section in Annex C of the IEA's WEO 2019¹³

[†]See <https://www.iea.org/data-and-statistics/data-tables?country=WORLD&energy=Balances&year=2018>, accessed 20 July 2021

sense. It is not known whether fossil fuels will continue to act as feedstock for such products in future, or whether replacement materials will be used. This issue is beyond the scope here. Note that the energy used for processing the feedstock in industrial applications *is* included in the industry sector of the model. If either petrochemicals or substitute non-petrochemicals are used, and their processing requires roughly equivalent amounts of energy, then the overall energy requirements will be the same across scenarios and do not need to be included here.

1.5.4 Bioenergy, solar thermal energy, marine energy and geothermal energy

We omit direct-use bioenergy (biomass, biofuels, biogases) because any significant contribution at a planetary scale is likely to have high environmental costs in terms of land use change, availability of water resources, as well as potential impacts on the costs of food¹⁴. We omit solar thermal energy and marine energy because of their high current costs and lack of progress historically. We omit geothermal energy because it is (currently) very location dependent. Solar thermal, marine and geothermal all have very low installed capacity levels and low recent growth rates, so that adoption at planetary scale on a timescale relevant for seriously contributing to emissions reduction appears unlikely. They may well improve and play valuable roles several decades from now, but more information would be needed to determine whether their costs are likely to drop sufficiently. Taken together, these only amount to about 4.0% of final energy, so their effect is relatively small.

1.6 Other energy system components

Sections 1.4 and 1.5 above refer to final energy components of the energy system, i.e. they cover the delivery of energy carriers for consumption by sectors. This section refers to the facilitating technologies required to allow the final energy flows to exist, i.e. transport, storage and distribution technologies (batteries, electrolyzers, P2X fuels, the power grid, fuel delivery networks, gas stations, electric vehicles, EV charging stations, and heat pumps). We also include in this section an explanation for why carbon capture and storage is not employed in this modelling framework despite its extensive use in many energy-economy models.

1.6.1 CCS

Carbon capture and storage (CCS) is a well established technology that has been operating commercially around the world for many decades¹⁵, since the successful introduction of CCS for enhanced oil recovery (EOR) in 1972¹⁶. There are several methods of EOR, one of which is CO₂ EOR, where emissions from nearby industry or power generators are captured and pumped in to an existing oil reservoir to aid the extraction of more oil. CCS costs are generally not well known, as this is sensitive commercial information for fossil fuel companies, which are the primary organizations active in the space at any scale. Where cost information has been made available, there is so far no evidence that costs have fallen^{17,18}, despite decades of investment in R&D programs, in both the public and private sectors¹⁹. There is also no evidence to suggest costs will likely fall in the future to levels allowing CCS to contribute seriously to emissions mitigation (without a high carbon price).

We omit CCS technologies from our model for three key reasons. First, their cost histories are not well known, so their wide commercial viability (i.e. beyond EOR) is unclear. Second, their production scale has always been relatively small. Third and most importantly, as described in more detail below, we are able to make simple worst-case assumptions that bound the impact of CCS costs on our results, so modelling them explicitly is unnecessary.

There are three potential major uses of CCS in the energy sector (recall that the scope of this model is the energy system only, so non-energy direct emissions are omitted). The first is CCS applied to fossil fuel power generation. The cost of this type of generation is bounded below by the cost of unabated generation, so modelling just the cost of unabated generation, as we do in this work, gives a lower bound on the cost of any system that includes unabated power generation *plus* CCS. This is useful in this study because in the results section (S.I. 7) we find that it is economically favourable to substitute even unabated fossil fuel power generation for renewables. Hence, the lower bound implies that retaining fossil fuel power generation and adding CCS would be an even more expensive scenario, compared to simply replacing it with renewables (as occurs in the Fast and Slow Transition scenarios).

The second major potential use of CCS is in capturing emissions from fossil fuels used in some hard-to-electrify applications in industry, in particular high temperature heat. Our transition scenarios assume that these fossil fuels are replaced by P2X fuels. This assumption results in an upper bound on the cost of eliminating emissions from industry in each scenario. It is possible that CCS solutions may be found that

offer a cheaper route to full decarbonisation of this sector than using P2X fuels, in which case clean energy transition scenarios could be even cheaper than the cost estimates given in this work.

The final major potential use of CCS is bioenergy with CCS (BECCS), which involves using bioenergy combustion for electricity generation, then capturing and storing the emissions. This results in negative emissions. In contrast to CCS used with fossil fuels, BECCS may be required in future given the uncertainties around residual non-energy emissions in sectors such as industry and agriculture (e.g. cement production or fertiliser use), or to lower greenhouse gas (GHG) in the atmosphere below the levels that we will arrive at once net-zero is achieved. With an estimated levelized cost of around 220 \$/MWh and a sequestration capacity of around 1.3 tonnes of CO₂ per MWh²⁰, BECCS could in future provide a competitive option for GHG removal. However, assessing such negative emissions strategies is beyond the scope of this research as it would require an emissions model for the entire global economy to assess comparative cost-effectiveness; an understanding of how the increased demand for bioenergy fuels might increase the cost of BECCS²¹; and an ability to cost the potentially adverse impacts of bioenergy fuels on the environment and food availability¹⁴.

1.6.2 Electricity technologies and networks

Electrification of energy is considered one of the three pillars of deep decarbonisation strategies, along with fuel switching and efficiency measures²². Even without the imperative to decarbonise, electricity is likely to become the world's dominant energy carrier at some point in the future, due to its favourable physical characteristics (versatility, controllability, cleanliness, lack of emissions or smell at point of use etc.). Although almost one billion people still lack access to reliable grid power²³, most countries have seen electricity consumption increase faster than primary energy demand since the turn of the century. The IEA estimates that between 1973 and 2017 primary energy consumption grew on average 1.9% per year while electricity generation grew at 3.3% per year²⁴.

Several electricity conversion, storage and generation technologies are included in our model: electrolyzers (S.I. 3.3, S.I. 6.13), two types of utility-scale grid battery (S.I. 3.1, S.I. 3.2, S.I. 6.12), and seven different types of generator. Suitable cost and production data exists for all these technologies, so we are able to make probabilistic cost forecasts and use these in our analysis.

The key facilitating technology on which all these components rely is the power grid. By “power grid” (or electricity networks) we mean transmission lines, distribution lines, substations, control systems, and any associated infrastructure, but *not* generators (or the fuels they require). We include power grid costs in the model. However, we use a slightly different method than for other technologies, because power grid costs and installations are harder to quantify, and suitable historical data is not available.

To keep the model as simple and as closely tied to available data as possible, we model the power grid in two steps, based on its two main functional characteristics. We call the first part power grid *flexibility*, meaning the ability of networks to manage and integrate intermittency and large fluctuations of variable renewable energy (VRE) sources. We include all flexibility related costs in the model by explicitly including appropriate quantities of grid batteries and electrolyzers to match solar and wind deployments. This allows us to make use of the detailed cost histories of these technologies.

The second power grid characteristic we model is its *size*, meaning its ability to handle a given amount of electricity generation (regardless of intermittency and variability considerations). To model this we assume that current annual investments in electricity networks scale in proportion to the amount of electricity generated in any given scenario, so that more net generation leads to larger annual investments. We also assume that the investment costs per unit of electricity generation remain constant into the future.

Using these two methods, we include in the model all costs associated with the very different electricity networks required across the different scenarios considered. For more details see S.I. 3.7.

1.6.3 P2X fuels

P2X fuels here refer to several fuels, including hydrogen, ammonia, methane, and methanol, among others²⁵. We assume that the starting point for all P2X fuels is electrolytic hydrogen from zero carbon electricity (“green hydrogen”). Currently almost all hydrogen is produced from fossil fuels, mostly from steam methane reforming (SMR), but also from coal gasification²⁶. Hydrogen produced from fossil fuels with unconstrained GHG emissions is known as “grey hydrogen” and accounts for almost all hydrogen currently produced. Fossil fuel hydrogen with CCS is known as “blue hydrogen”, though this is rarely produced due to the extra costs incurred by CCS.

We do not include grey or blue hydrogen in our model for several reasons. First, grey hydrogen has very high GHG emissions so is not sustainable in the long term. Second, blue hydrogen is dependent on CCS technologies, which as discussed earlier are expensive, have no history of cost declines or significant technological progress, and are unproven at scale. Third, the feedstock for both grey and blue hydrogen are fossil fuels, whose prices have been approximately constant for around a century (see S.I. 6), so even if blue hydrogen is brought to scale, it is unlikely that there will be significant reductions in this major component of the cost. Finally, both steam reforming and gasification are mature technologies, so there is little reason to expect any significant reduction in their capital cost components either.

We do not attempt to model the final mix of P2X fuels as this does not dramatically affect the resulting costs. We only model power-to-hydrogen, which represents a majority of the cost for all common P2X fuels. For example, producing ammonia via hydrogen currently costs about 15-25% more than the cost of producing hydrogen alone²⁶, while the premium for synthetic methane is around 30%²⁵. Importantly though, these production premiums are offset by lower transport and storage costs, so total costs are roughly equivalent. The largest cost components by far are the electricity and the electrolyzer technology, which we model. We discuss these issues in more detail in S.I. 3.

1.6.4 Fuels infrastructure

Infrastructure costs vary for different P2X fuels, just as they do for fossil fuels, and it is assumed that different kinds of fuels will be available for different applications (i.e. a range of gases and liquids, stored at various temperatures and pressures etc.).

Hydrogen infrastructure currently costs around 2-3 times as much as oil and gas infrastructure per unit of energy (see S.I. 3). This is partly because hydrogen gas has very low energy density (around 1/3 that of natural gas), so more compression is required during transmission and storage. But it is also partly due to the fact that current networks are small and have not realized economies of scale. There are currently roughly 5,000 km of hydrogen pipelines around the world vs. 3 million km of natural gas transmission pipelines²⁶. Recent trials and commercial projects have demonstrated that much natural gas infrastructure can be repurposed for hydrogen delivery, and costs are expected to fall with scale^{27,28}. Infrastructure costs are generally not believed to be a significant impediment to deployment (it is the cost of the hydrogen itself that has so far prevented more widespread use, currently around 3-4 times the cost of natural gas).

Due to its higher density, ammonia is generally cheaper and more convenient to store and distribute than hydrogen²⁶. There is already a global scale ammonia production and distribution industry, mostly related to fertiliser production, consisting of pipelines and ships similar to those used in the fossil fuel industry (chemical tankers and liquid petroleum gas tankers), and there are currently over 7000 km of ammonia pipelines.

Due to a lack of historical time series data on infrastructure costs for pipelines and storage, future infrastructure costs are difficult to estimate. We thus make the pessimistic assumption that there will be no cost improvements, i.e. that future infrastructure costs for hydrogen and ammonia per unit of useful energy delivered will be equal to current costs. In the Fast Transition scenario the size of the required P2X industry, in terms of annual energy supplied, is around 1/3 - 1/2 the size of the combined oil and gas industries in alternative scenarios (depending on the year). This is because most of the energy system is electrified, so fuel requirements are much lower. The remaining fuel-based energy distribution and storage network in this scenario is consolidated in hubs around industrial uses, shipping and aviation. This results in a down-scaling of the network by a factor of 2-3. Thus the unit cost difference cancels out: total infrastructure costs for P2X will be less than or equal to total infrastructure costs for fossil fuels. Hence total P2X infrastructure costs in the Fast Transition scenario will likely be lower than those required to maintain, upgrade and expand the fossil fuel infrastructure in any other scenario. We pessimistically assume no cost decreases as the P2X energy distribution and storage network grows to scale. We discuss these issues in more detail in S.I. 3.

To put this in perspective, a kilometer of oil pipeline costs roughly a million dollars, so the entire global pipeline system is worth about 3 trillion dollars. Even if we had to rebuild a global ammonia pipeline system from scratch, it would cost about 3 trillion dollars, which is a large number but is still smaller than the 12 trillion of expected savings that we estimate for the Fast Transition scenario. More realistically, depreciation rates for oil pipelines vary from roughly 3% to 14%, so the entire fossil fuel pipeline system would effectively be replaced over a timescale ranging from 7 to 30 years. Therefore the sunk costs are likely to be small. In any case, we expect that for a fast transition a substantial fraction of pipelines will be retrofitted for ammonia, hydrogen, or other P2X fuels.

1.6.5 Electrification of transport

Our model is primarily focused on the supply-side, i.e. on energy technologies such as electricity generators, batteries, fossil fuels, and P2X fuels. For hard-to-electrify sectors, such as heavy-duty transport and high-temperature heat in industrial processes, scenarios are constructed to ensure that in any high renewables scenario, sufficient quantities of P2X fuel are available to avoid significant disruption to demand-side operations. However, the one really significant demand-side technology change that will be required for any high renewables scenario is the widespread adoption of light-duty EVs.

To model the transition from internal combustion engine vehicles to electric vehicles in the light-duty vehicle market segment, we simply neglect the costs of capital stock replacement. This is because *i*) EVs are much simpler and have far fewer moving parts than ICEVs, and so are likely to be cheaper in most market segments by the mid-2020s²⁹, and *ii*) the entire stock of light-duty vehicles naturally turns over within a few decades in all scenarios. Thus, by excluding this component from the model, our net system cost estimates provide worst-case cost estimates associated with transitioning away from a No Transition scenario: in the worst case it would cost as much to replace ICEVs with EVs as it would to replace them with ICEVs, and the cost of stock replacement would be equivalent across all scenarios. The Fast and Slow Transition scenario cost estimates are therefore overestimates, relative to No Transition, since the likely capital *savings* from EV capital costs are omitted. This simplification allows us to focus on modelling the cost of providing the energy, i.e. the cost of gasoline vs. electricity, which is the factor that dominates the cost differential between scenarios.

Regarding EV charging infrastructure, there are two components to consider: grid reinforcements to accommodate higher levels of power and net energy transfer; and charging stations. The cost of grid upgrades in our model are already accounted for in our electricity infrastructure assumptions (S.I. 3.7), with higher electrification scenarios having correspondingly higher grid investments. Then, regarding charging stations, since the electric grid is already ubiquitous in the vast majority of locations where EVs would be used, and EVs can, if required, be plugged in to standard low power outlets, charging station costs can in theory be very low. Indeed, most EV charging currently takes place either at home or at work³⁰. However, some users will need faster charging solutions that require specialised infrastructure. We do not account for these potential extra costs, though we note that they have been falling rapidly and are expected to continue to do so³¹. Furthermore, it is important to remember that fuelling infrastructure for fossil fuels is costly to build and maintain, and has relatively short lifetimes and substantial turnover, e.g. a typical depreciation time for a gas station is 15 years³², so new rapid-charging electrical infrastructure can be built on the timescale required as old infrastructure is retired, with a likely net saving.

1.6.6 Other demand-side capital costs

Most demand-side components of the energy system will need to be replaced by 2050 in all scenarios. For example, virtually all heating and cooling systems and end-use appliances will need to be replaced between now and 2050 in all scenarios. (Only a small fraction of ships and planes will have lifetimes over 30 years, and we account for hard-to-electrify applications such as these by providing sufficient quantities of fuel, either fossil-based or P2X depending on the scenario.) There are no significant differences in investment costs between appliances that run on different energy carriers, and so for simplicity we omit all such capital stock replacement costs from the model entirely. The one notable exception to end-use technologies with lifetimes longer than 30 years are buildings. However, we do not assume any energy efficiency improvements, such as improved insulation – we only assume that fossil fuel based heating systems are replaced by electricity. In our spirit of making conservative estimates, we assume that this is resistance-based heating, which is highly inefficient but can be done with very cheap infrastructure (i.e. we assume that in the buildings sector electricity is converted to useful energy with a conversion efficiency of 1, see S.I. 1.8). A typical replacement time for a gas heater is around 10 years, making it possible to replace all heating systems within the timescale of the Fast Transition at no net cost. This is a conservative estimate – installation of better heating solutions such as heat pumps are likely to substantially improve the efficiency and lower the overall cost of the transition relative to what we have estimated here.

1.7 Units and justification for the use of LCOE

In order to keep our model as simple as possible, we deal almost entirely in terms of what is sometimes loosely called “energy” (i.e. annual electricity generation, annual fuel supply, or installed energy storage capacity), as opposed to what is loosely called “power” (i.e. installed capacities of electricity generators, or charge/discharge rates of installed energy storage technologies). For electricity generation technologies, we

use energy generated (MWh) and LCOEs (\$/MWh), as opposed to installed capacities (kW) and capacity costs (\$/kW). Most energy system models describe electricity supply in terms of installed generation capacities, then have to model utilisation rates in order to convert power to energy; our method avoids this. Direct-use fossil fuels are usually modelled in terms of energy content (i.e. MJ or MTOE or MMBtu); we also use this method, and hence most technologies in our model are described in terms of energy and cost per unit energy (see Table S5).

The technologies included in our model are usually described in units specific to each technology, which are not common across technologies (for example, see Table S24). To make all quantities and costs comparable in computations, we have converted almost all production data to units of exajoules (EJ), and cost data to units of dollars per gigajoule (\$/GJ; 1 \$/GJ = 1 billion\$/EJ). The only exceptions to this energy-only perspective are electrolyzers, which are described in terms of installed capacity (these are usually described in terms of kW and \$/kW but we convert them to EJ hr⁻¹ and \$/GJ hr⁻¹), and electricity network investments, which are given in dollars per year.

There is no single natural metric for measuring “the cost” of any energy technology (or even any technology in general). This is because each technology has many different properties that together determine how it is used as part of a wider, evolving technology system³³. In this work we use LCOE as the cost metric for all electricity generation technologies, as it provides a good all-in approximation of the cost of providing each unit of energy.

LCOEs are determined by a calculation involving several key technology parameters, including capital costs, fuel costs, operation and maintenance costs, capacity factors, project lifetime, interest rates etc³⁴. While many energy system models calculate LCOEs by forecasting each of these factors separately, it makes more sense in our model to just forecast LCOEs directly based on observed data. This is because, first, our cost forecasting method is calibrated using observed historical data, and the longer the data series used for calibration the better (typically the forecast uncertainty decreases as more data is used for calibration). Attempting to collect and use data on each input factor separately will in practice limit the length of the historical record available for making forecasts, and hence result in unrealistically wide forecast error bars. Second, and more importantly, the statistical validity of probabilistic forecasts produced by aggregating forecasts of technology sub-components in this way has not been tested, and hence such forecasts do not necessarily have the same statistical properties as the corresponding single series forecasts do. Technology sub-components are likely to be highly correlated if, for example, each sub-component’s progress depends on a universal background factor, such as the annual level of technology-specific R&D expenditures. Although this has not been tested rigorously, it would be expected that failing to account for such correlations, and instead naively aggregating forecasts of correlated sub-components in this way would result in higher forecast uncertainty than obtained using the single series method, and that this would simply be an artifact of the aggregation procedure, rather than a reflection of any “true” technological uncertainty.*

Now, since observed LCOE data is used to calibrate our forecasting model and generate LCOE forecasts, this means that the historical variation in each separate cost component of LCOE is implicitly fed forward in to the forecasts, and is reflected in the width of the forecast error at the LCOE level. For example, PV module efficiencies have risen from around 12.7% in 2002 to 19.5% in 2020^{35,36}, and this improvement is captured in our forecasts. Similarly, PV capacity factors have gradually increased, interest rates decreased, and lifetimes increased. While we can not tease out from our LCOE forecasts the exact contribution of each component, it is not necessary to do this in this study, it is enough to understand that the forecasts encompass the full range of uncertainty implied by the observed historical data, and reflect the expected combined changes in all LCOE components.

One particular concern, however, relates to the possibility of average capacity factors decreasing as the energy transition proceeds. As solar and wind deployment increase, it is possible that significant load-shedding may occur if more renewable capacity is installed than is instantaneously required and there is nowhere for the electricity to go. If this became routine practice worldwide, then capacity factors may fall, which would – if all other LCOE input variables remained static – lead to a rise in average LCOEs. While this may initially

*Note that this criticism could be made at the system level for our probabilistic system cost forecasts – how sure are we that separate energy technology forecasts can be combined to generate a reliable total system cost forecast? While we admit that aggregating forecasts in this way has not been “tested”, first, there appears to be no better existing method to generate such system level probabilistic forecasts. Second, we do not claim that our probabilistic system cost forecasts themselves are empirically validated, just that the single technology forecasts on which they are based are. Finally, the technologies included in our model are generally very distinct in character (this is one reason why we do not include onshore and offshore wind as separate technologies, for example), and hence are far less likely to be as closely correlated as sub-components of any specific technology, so the resulting forecast errors are not expected to simply be an artifact of the aggregation process.

appear problematic for our method, it is not. The primary reason why it is not a problem is that in our model, high VRE penetration is always accompanied by high levels of storage and transmission capacity, which are expected to absorb most of the over-generation, most of the time. Indeed, the model is designed around this behaviour. Since in the Fast Transition scenario annual electricity generation reaches around 140% of annual final electricity demand, the system will need to be configured around this requirement. Detailed power system planning is well beyond the scope of this study, highlighting the complementary nature of our approach with both more detailed integrated assessment models, and high resolution, system-coupled, power system analysis models (e.g. Brown et al. ³⁷). In addition, policy tools are available for addressing the question of the falling market value of solar and wind electricity with increasing penetration ³⁸.

Another reason why the possibility of LCOEs increasing is not a problem in our model (whether due to decreasing capacity factors or any other reason) is that our cost forecast probability distributions are log-normal, and therefore already include a long tail of potentially high costs. That is to say, for any given technology the forecast inherently includes a non-zero probability that costs will rise. The forecast models are calibrated on historical data, which inevitably contains some year-to-year cost increases, and so the forecasts automatically reflect the fact that such increases may occur again in future, for a variety of reasons. For PV, for example, module costs increased for several consecutive years in the mid-2000s, leading to LCOE increases. So although PV has not witnessed a sustained decrease in capacity factors historically, it has seen sustained module cost increases. Hence our PV LCOE forecasts do account for the possibility that any of the LCOE input variables may reverse progress, leading to LCOE increases in future. Although this reasoning is sound, it could be claimed that perhaps it does not account for the kind of changes in LCOE input variables that might result from a structural shift in the energy system, as embodied in the Fast Transition scenario. To counter this, first note that even if capacity factors decreased, LCOEs may well continue to decrease if sufficient progress is maintained in other key parameters, such as system costs, module efficiency and lifetimes (all of which appear likely). Second, the forecasting model has been tested on data for over fifty technologies from over many different decades, so many of these are likely to have experienced such structural shifts, of one sort or another, in their particular domain, and so there is no reason to believe the forecasts are not robust to this eventuality.

To address what is essentially the same issue, but from a slightly different perspective, one feature of the LCOE metric is that it does not account for electricity system integration costs (i.e. the value of dispatchability, inertia, frequency response etc. to the wider system). This is to be expected, as no single cost metric can account for all costs and benefits associated with using a technology as part of a wider system – there will always be many external factors affecting costs, and these will always depend on how the system developed historically (i.e. the extent to which existing infrastructure has created technology lock-in). Electricity system costs are sometimes quantified in terms of an extra component added to the LCOE of each generation technology (e.g. VALCOE, the value-adjusted levelised cost of electricity, see IEA ¹²). However, calculating this component requires very detailed modelling (e.g. OECD & NEA ³⁹), which necessarily involves many more technical assumptions and layers of complexity. Instead, as described above and in S.I. 3, we have ensured that all system costs are accounted for by including (overly) generous quantities of storage capacity, dispatchable generation and grid expansion, and adding their costs separately to the total system cost calculation.

It would of course be possible to build a model using a technology cost metric other than LCOE, such as investment costs, but this would require many more assumptions of other technology parameters such as capacity factor, as discussed above. For a discussion of the choices of variable available when applying Wright's law in the context of energy technologies, see Trancik et al. ⁴⁰.

1.8 End-use conversion efficiencies

Sectors of the economy convert final energy into useful energy to provide energy services. Each energy carrier has an average final-to-useful energy conversion efficiency in each sector, based on the types of processes for which it is used; we denote this by $\eta_{carrier}^{sector}$. We assume that this average value is broadly representative of the carrier's conversion efficiency in all applications throughout the sector (for example we assume that most of the final energy in gas, when used by industry, is converted to useful energy with an efficiency of roughly $\eta_{gas}^{industry}$). Evidence for this assumption, and the specific values we use, are given in De Stercke ⁴¹. It is beyond the scope here to predict how these efficiencies may change in the future, so we assume that they are constant and remain fixed for the duration of model time. Sector-specific carrier conversion efficiencies used in the model are shown in Table S4. Also shown are recent final energy values and the useful energy values implied

by the conversion efficiencies.* †

Sector	Carrier	Final → Useful efficiency, $\eta_{carrier}^{sector}$	2018 final energy, Mtoe	2018 final energy, EJ	2018 useful energy, EJ
Transport	Oil	0.25	2629	110.1	27.5
	Electricity	0.8	32	1.34	1.07
	P2X fuels	0.5	-	2.95×10^{-4}	1.48×10^{-4}
	All		2661	111.4	28.6
Industry	Oil	0.6	298	12.5	7.5
	Coal	0.6	795	33.3	20.0
	Gas	0.6	646	27.0	16.2
	Electricity	0.8	803	33.6	26.9
	P2X fuels	0.6	-	2.95×10^{-4}	1.77×10^{-4}
	All		2542	106.4	70.6
Buildings	Oil	0.7	330	13.8	9.7
	Coal	0.6	125	5.23	3.14
	Gas	0.6	700	29.3	17.6
	Electricity	1.0	1011	42.3	42.3
	P2X fuels	0.6	-	2.95×10^{-4}	1.77×10^{-4}
	All		2166	90.6	72.7
Energy	Oil	0.6	331	13.9	8.3
	Coal	0.6	410	17.1	10.3
	Gas	0.6	394	16.5	9.9
	Electricity	1.0	371	15.5	15.5
	All		1506	63.0	44

Table S4: Carrier and sector specific conversion efficiency assumptions. $\eta_{carrier}^{sector}$ is the final to useful energy conversion efficiency of the given carrier in the given sector. These values are based on the data and analysis presented in De Stercke (2014). Final energy data for 2018 is shown (from the World Energy Outlook¹³), along with the useful energy values implied by the assumed conversion efficiencies.

Providing energy services by using energy carriers with higher conversion efficiencies can be one of the easiest ways to reduce final energy consumption. For example, to provide a given level of transport service, a transport sector that uses oil as its main energy carrier (i.e. composed of mostly ICEVs) requires approximately three times as much final energy as a transport sector that uses electricity as its main energy carrier (i.e. composed of mostly EVs). Thus moving to a fleet of EVs drastically reduces the transport sector's final energy requirements for a fixed service level. This carrier substitution principle holds in each sector, with greater electrification leading to lower energy wastage.

Of particular note is the conversion efficiency value of 1 chosen for the use of electricity in buildings. This was based on historical data, and represents the case where any electricity-based heating services are provided with resistance heaters only (since these have efficiency 1). However, this may change considerably in a future where heating is mostly provided by heat pumps. These devices use ambient energy drawn from the environment, and have a coefficient of performance (COP) (i.e. conversion efficiency) in the range 2-5, depending on conditions⁴². If heat pumps become widespread, as appears likely, the average conversion efficiency of electricity in buildings could conceivably rise to $\eta_{electricity}^{buildings} \approx 2$ or higher, which would significantly reduce the quantity of final energy required by the sector. Thus our value of 1 here is conservative.

*Regarding the energy sector, note that the sector total (1506 Mtoe) and the electricity component (371 Mtoe) are found in WEO 2019 (see Table S3). For oil, coal, and gas though, the final energy values are calculated by splitting the remaining 1135 Mtoe between these three carriers in proportion to their energy sector consumption as shown in the 2018 IEA energy balance table available online at <https://www.iea.org/data-and-statistics/data-tables?country=WORLD&energy=Balances&year=2018>, accessed 20 July 2021.

†While most of these final energy values are easily obtained from the IEA's World Energy Outlook, this is not the case for P2X fuels, so we estimate these values separately. Our model only considers P2X fuels produced from electrolysis, and most P2X fuels are currently not produced this way, so we must take care not to overestimate these values. Approximately 80 MW of hydrogen electrolyzers (of the appropriate type) currently exist. Assuming a utilisation rate of 50% and conversion efficiency of 70% yields a current annual quantity of P2X fuels of 8.86×10^{-4} EJ/yr. For simplicity we suppose that this is split equally among the three end-use sectors, giving values of 2.95×10^{-4} EJ. These quantities are so tiny that the specific values used are unimportant, all that matters is that the orders of magnitude are approximately correct. See S.I. 6.13 for further details.

1.9 Estimating total system costs

As established in Table S2, currently about 81% of final energy in the energy system is covered by direct-use oil, coal and gas, plus electricity. Therefore, by forecasting the costs of just these four energy carriers, plus P2X fuels as a potential substitute for fossil fuels, we can obtain a first-order approximation of the future total system cost. To form a more accurate estimate there are two extra types of cost that we must consider.

First, there is the remaining 19% of final energy not covered explicitly by model components. Second, there are the many additional technologies required to facilitate the transport, distribution, storage and use of these energy carriers. Our model covers almost all of these extra energy system components, either by including extra technologies in the model (e.g. batteries, electrolyzers, electricity networks), or by making simple worst-case assumptions regarding the costs of omitted components, and how they will vary in any transition away from the current fossil-fuel-based system. This means that as well as forecasting total costs of different scenarios in our model system, in absolute terms, we also forecast *relative* costs of different scenarios, which apply to both our model system and the full, real-world system.

Consider the 19% of final energy not covered explicitly by model components. Within this block are bioenergy, biofuels, other renewables, and traditional biomass. As shown in Table S2, these sum to 45.5 EJ of final energy, which is just under 10% of final energy. These minor energy carriers are all “renewable” technologies already (perhaps to varying extents depending on how they are sourced). If we assume that their costs and production will evolve independently of the major energy carriers included in the model, and independently of which scenario is implemented, then they will contribute the same costs to all scenarios. Our relative scenario cost estimates will then be valid whether or not they are included. The model therefore covers 90% of global final energy, excluding these renewable energy carriers (and petrochemical feedstock). The remaining omitted components are heat (used as an energy carrier), a category called “other fuels”, and another category called “unspecified”, that together constitute the last 10% of final energy. The latter two are not specified in detail in the IEA data tables used to construct the model, so we are unable to include them.

Next consider additional energy infrastructure technologies. These were presented in detail in the discussion of the various energy system components above (S.I. 1.6), but the largest components are summarised as follows:

- *Transport, distribution and storage of electricity.* Large quantities of utility-scale grid batteries, electrolyzers, and P2X fuels are included in the Fast and Slow Transition scenarios to integrate VRE sources; these costs are included explicitly in the scenario NPC forecasts. The costs of increasing power grid capacity to handle more electricity generation are also included explicitly in the model, see S.I. 1.6.2 for further details.
- *Transport, distribution and storage of fuels.* These facilities can be provided in Fast and Slow Transitions at no extra cost relative to No Transition because although unit costs may be slightly higher for P2X fuels than for fossil fuels, a corresponding down-scaling of the fuels sector occurs, as the economy undergoes large-scale electrification. Therefore, despite not being modelled explicitly, these costs are implicitly included in the scenario relative NPC forecasts. See S.I. 1.6.4.
- *Electric vehicles.* These can be provided in Fast and Slow Transitions at reduced cost relative to vehicle expenditures in No Transition, because EV capital costs will be lower than ICEV capital costs. Thus, although we do not model vehicle capital stock explicitly (and hence omit the likely EV savings entirely), these costs are implicitly included in the scenario relative NPC forecasts. See S.I. 1.6.5.

As a result of these assumptions, we are able to forecast relative system costs in different scenarios by simply summing up the costs of the 14 technologies shown in Figure S1, in each scenario. These are summarised in Table S5.

Energy system component	Cost included
Oil (direct use)	Cost per unit energy
Coal (direct use)	Cost per unit energy
Gas (direct use)	Cost per unit energy
Coal electricity	LCOE
Gas electricity	LCOE
Nuclear electricity	LCOE
Hydroelectricity	LCOE
Biopower electricity	LCOE
Wind electricity	LCOE
Solar PV electricity	LCOE
Daily batteries	Annuitized capital cost
Multi-day storage	Annuitized capital cost
Electrolyzers	Annuitized capital cost
Electricity networks	Annual investment cost

Table S5: The 14 technologies included in our model, the costs of which are used to form scenario cost estimates relative to No Transition. (All levelized and annuitized costs remain fixed for the duration of each technology’s lifespan, initialised at its year of installation – see S.I. 3.8 and S.I. 5.3.2.)

Table S6 summarises all energy system components that are modelled in our total cost estimates, and those that are not.

Energy system component	Represent explicitly in model?	Represent implicitly in model?	Notes
Fossil fuel production, conversion and use (S.I. 1.4, S.I. 1.4.2, S.I. 6.2-6.6)	Yes		Direct-use fossil fuels are included in the model, in all sectors. This includes the energy sector itself, which accounts for conversions to high quality end-use fuels (e.g. from crude oil to gasoline). Power generation fossil fuels are included via the quantity of electricity they generate, and their respective LCOEs
Petrochemical feedstock (S.I. 1.5.3)	No	No	We exclude this as it does not provide energy services.
Electricity generators (S.I. 1.4, S.I. 1.5, S.I. 1.6.1, S.I. 6.5-6.11)	Yes		All major technologies are included. CCS, marine, tidal, and geothermal electricity are excluded as they have low installed capacity, low recent growth rates, and do not have promising cost trends.
Electricity storage (S.I. 1.6.2, S.I. 6.12)	Yes		Daily grid batteries, multi-day grid batteries and P2X fuels for power grid backup are included.
P2X production (S.I. 1.6.3, S.I. 6.13)	Yes		The cost of hydrogen production is included in the model via electrolyzer system costs and electricity costs. Other P2X fuels may be produced as well, but any additional costs are not included, though these are relatively small and would likely be offset by lower storage and transportation costs.
Electricity transmission and distribution network (S.I. 1.6.2)	Yes		Power grid flexibility costs are included via the inclusion of numerous storage technologies. Power grid capacity expansion requirements are included via annual network investments that scale with electricity generation in each scenario.
Non-storage ancillary grid services *	No	No	The costs of ancillary services beyond storage are likely to be very small in the Fast Transition, due to digitisation, demand management and improved control systems required to integrate VREs and storage.
Liquid and gas fuel storage and distribution network (S.I. 1.6.4)	No	Yes	Total infrastructure costs for P2X are assumed to be less than or equal to total infrastructure costs for fossil fuels. Hydrogen infrastructure currently costs around 2-3 times as much as oil and gas infrastructure per unit of energy but in the Fast Transition scenario the size of the required P2X industry is 2-3 times smaller than the combined oil and gas industries in No Transition
Vehicles (S.I. 1.6.5)	No	Yes	Virtually all vehicles must be replaced within 20 years, so transitioning from ICEVs to EVs would likely save money, due to their expected lower capital costs. Hard-to-electrify vehicles (e.g. heavy-duty trucking, ships, airplanes) are provided with ample P2X fuels in transition scenarios. However, it is likely that cheaper options will arise (e.g. improved battery technologies), reducing the costs of transition scenarios further.
Heating assets (S.I. 1.6.6)	No	Yes	Virtually all heating must be replaced by 2050 in all scenarios. In high electrification scenarios, enough electricity is produced to provide heating services using only cheap, inefficient resistance heating, so that investment costs are equivalent across all scenarios. However, if heat pumps are used instead then their higher efficiencies will decrease electricity requirements significantly, and reduce the costs of transition scenarios further.

* This refers to a variety of technical services required to maintain stable power flow in the electrical grid, e.g. scheduling, dispatch, operating reserves, and frequency and voltage control.

Table S6: Table outlining which energy system components are included in our scenario relative cost forecasts.

The scenarios are designed specifically so that the results reflect upper bound cost estimates for the Fast and Slow Transition scenarios relative to No Transition. There are many conservative assumptions made and

discussed in various sections of the manuscript, but we summarise them all in one place here for clarity and to emphasise the conservative nature of our scenario cost comparison. We list here instances where we assume the minimum possible cost for the No Transition scenario, or the maximum possible cost for the Fast Transition scenario.

- Oil and gas both exhibit weak long-term increasing price trends, as demonstrated by over a century's worth of data. However, in our model we assume that oil and gas costs continue at roughly their current levels into the future. This is a conservative choice, particularly in the No Transition scenario since demand increases significantly. The model therefore likely underestimates fossil fuel expenditures somewhat in the Slow Transition scenario, and significantly in the No Transition scenario.
- Our systematic use of technologies with comprehensive cost or price histories means that in many cases we neglect solutions that are promising avenues of cost reduction, such as demand-side management, heat pumps, and end-use efficiency improvements (e.g. improved building insulation). Including such solutions would lower costs for transition scenarios relative to No Transition (e.g. since less storage would be required).
- We use a conservative estimate on the maximum extent to which industry can be electrified. We assume that only 75% of useful energy in industry can be electrified, although recent analyses suggest 99% is achievable in Europe⁴³, and therefore a similarly high fraction should be possible globally. If a higher value such as 90% was used, then industry would not have to be supplied with so much relatively expensive P2X fuel, and cheap electricity could be used instead. Industry energy costs would be substantially lowered in our transition scenarios, as would the total system cost estimates.
- The conversion efficiency value chosen for the use of electricity in buildings (i.e. 1) is based on historical data, but may improve considerably in future if heating is mostly provided by heat pumps. Heat pumps have coefficient of performance⁴² in the range 2-5, so if they are used widely then the average conversion efficiency of electricity in buildings could rise to 2 or higher, which would significantly reduce the quantity of final energy required, and reduce costs. The assumption we use may therefore favour the No Transition scenario, since the use of simple resistance heaters in the Fast and Slow Transition leaves significant potential to cut costs by the use of heat pumps.
- We assume that transport and grid batteries are completely different stocks of the same technology. The quantities provided in scenarios assume no overlap through mechanisms such as EVs providing vehicle-to-grid storage services to the grid. This is perhaps a very conservative (and expensive) assumption that favours the No Transition scenario, as in reality there may be a great deal of overlap in their uses, meaning that the total quantity of storage required could be much smaller than we assume in the Fast and Slow Transition scenarios⁴⁴.
- We do not explicitly model the potentially large capital cost savings that would likely be generated by transitioning from more expensive ICEVs to cheaper EVs in the Fast and Slow Transition scenarios.
- Annual deployment growth rates for solar, wind, batteries and electrolyzers over the last twenty years have been, in many years, much higher than the rates that we use in the Fast and Slow transition scenarios. For example, in around 5 of the last 20 years, solar PV's annual growth rate was above 50%, but we use an initial growth rate of 32% in the Fast Transition scenario. It may therefore be possible to construct faster transition scenarios while still being broadly consistent with recent historical trends. In keeping with our conservative approach though we choose growth trends in the Fast Transition that are conservatively low.
- We use PEM electrolyzer technology for our electrolyzer cost forecasts. PEM has the high flexibility required to absorb VRE sources, is safe and reliable, and has good lifetime and efficiency characteristics, so it is a promising candidate. However, the lowest current cost estimate of AEL electrolyzers is already 80% cheaper than the initial PEM cost estimate we use, so we are being highly conservative in our cost estimates for electrolyzers.
- In our model we assume a 12 year lifespan for daily grid batteries, and 10 years for EV batteries. However, they will both continue to function at lower performance for much longer than this, perhaps another 10 years, so we anticipate very large quantities of second-hand batteries also being available for grid applications. Thus both our storage cost and system reliability estimates for the Fast and Slow Transition scenarios are conservative.

- P2X infrastructure costs are likely to decrease with scale. Once the P2X network is scaled-up to the size envisaged in the Fast and Slow Transition scenarios, it is likely that overall storage and distribution costs per unit of energy delivered would be broadly the same as those for fossil fuels delivered through the existing network, in both cases accounting for only a small fraction of delivered energy costs. We conservatively assume that P2X infrastructure costs remain at their current levels, 2-3 higher than fossil fuel infrastructure costs, so that a P2X distribution network of a given size costs the same as an oil and gas distribution network 2-3 times the size. We thus neglect the potential cost savings associated with scaling down the very large storage and distribution networks currently used for fossil fuels.
- We assume that the multi-week lulls in variable renewable energy outputs possible in some regions of the world are broadly representative of all global regions. In the Fast and Slow Transition scenarios we therefore specify that there must be enough P2X fuel stored to cover one month's worth of final electricity each year, everywhere. Reduced output for many weeks at a time is not likely in all regions, so providing enough P2X fuel to cover global final electricity for a month may be excessively conservative. (This is of course a critical issue though, which must be addressed correctly in each location based on its resources and connectivity.)
- It is possible that the solution we employ in this model to manage VRE sources (i.e. bulk energy storage), is one of the costliest methods available⁴⁵. Thus, our approach should be viewed as providing a conservative upper bound on the cost of integrating large shares of VRE sources in the energy system. Once again this assumption increases the assumed costs of the Slow and Fast Transition scenarios in comparison with the No Transition scenario.
- In calculating the power grid investments required to accommodate high levels of electrification in the Fast and Slow Transition scenarios, we have assumed that investment scales with total electricity generation, rather than the lesser amount of final electricity. This likely overestimates the costs because historically the majority of grid investments have been required by the dense grid networks that deliver electricity to end users, which are associated more closely with final electricity than total generation. Several other conservative assumptions are laid out in S.I. 3.7.
- Coal-fired electricity LCOEs currently range from around 50 \$/MWh to well over 100 \$/MWh. The data series we use to calibrate and initialise our coal power cost forecasts is a series at the very low end of costs. The coal LCOEs we use in the model therefore represent the cheapest, most polluting coal power stations. Since we focus only on minimum possible system costs here, without considering environmental externalities, we use this forecast, with the caveat that it relies on society continuing to accept very high levels of air pollution and carbon emissions; this is a very conservative assumption.
- We use US LCOE values to calibrate and initialise our cost forecasts for gas-fired electricity. But US costs are very low compared to most other locations because local gas supply is abundant and transportation by pipelines is cheap. Gas LCOEs in Europe and Japan are much higher, reflecting higher transportation costs. Hence the model likely underestimates global gas-fired electricity costs, which favours No Transition since this scenario includes far more gas-fired electricity than the others do.

In summary, the total costs of the energy system components not explicitly included in our total system cost estimate are very likely to be greater for the No Transition scenario than the Slow or Fast Transition scenarios. By purposefully excluding these components (for model simplicity and comprehensibility), in addition to making consistently conservative assumptions regarding parameter choices, model calibration, and model initialisation, we are most likely substantially overestimating Fast and Slow Transition scenario costs, relative to No Transition scenario costs.

2 Scenario construction

We follow Grubler et al.⁴⁶ in focusing on the provision of energy services. This perspective highlights the importance of the typical efficiencies with which different energy carriers provide energy services, and the large efficiency savings available by switching to more efficient energy carriers. Our strategy is to construct scenarios that meet the needs of a growing global economy (and perhaps exceed them by far). We aim for approximately correct quantities of all the various components from which an energy system may be constructed: electricity, direct-use fuels, and energy storage, transport and conversion technologies. The process is described in detail in the rest of this section, but we first provide a high-level summary. To construct *a single scenario*, the steps are:

1. Specify the quantity of useful energy consumed in each sector (excluding the fossil fuel components of useful energy in the energy sector – these are dealt with in step 7). These useful energy quantities are defined exogenously, and are identical across all scenarios.
2. Specify the quantities of energy carriers used in each sector (again, excluding energy sector fossil fuel energy carriers).
3. Specify the proportion of non-dispatchable generation in the electricity mix at the end of the model time horizon (i.e. in 2100).
4. Increase P2X fuel production (above the level given in point 2) if required to account for VRE intermittency (i.e. for grid backup). This depends on the VRE proportion specified in point 3.
5. Increase total electricity generation (above the level given in point 2) if required to account for electrolytic production of P2X fuels. This depends on both the quantity of P2X fuels in end-use sectors (point 2) and the quantity required for grid backup (point 4).
6. Specify the quantity of electricity produced by each electricity generation technology.
7. Calculate the quantities of fossil fuel energy carriers required by the energy sector. This is a function of the quantities of fossil fuels used in end-use sectors (both in electricity and as direct-use fuels). This depends on points 2 and 6 above.
8. Calculate the quantity of daily-cycling batteries required. This is a function of the quantity of non-dispatchable generation (point 3), and the level of electrification of the transport sector (point 2).
9. Calculate the quantity of multi-day storage batteries required. This is a function of the quantity of non-dispatchable generation (point 3).
10. Calculate the quantity of electrolyzers required for P2X fuel production. This depends on point 4.

We use a time horizon of 2070 for our main results, but also present robustness checks on different time horizons up to 2090 in S.I. 7.2 and S.I. 7.4. Scenarios must therefore exist to at least 2090, and in fact we construct scenarios out to 2100. However, since all transition dynamics have occurred by 2070, we only show scenarios in detail up to this point; additional detail beyond this time is unnecessary as it has no bearing on our results.

2.1 Growth rate model

Scenarios are constructed by growing or contracting certain model components at different rates over time. These growth-rate-determined components are the primary drivers of scenarios, and are used to determine various other dependent components. To implement these growth processes a common framework is used: growth rates follow logistic growth curves (i.e. S-curves). Each such model component begins in an initial state matching its current production level, and approximately matching its historical growth rate.* Then over time it gradually transitions to a new configuration, ending up in a new steady state defined by its long-term growth rate.

*There are a few exceptions, for example in order to implement the No Transition scenario, the growth rates of solar PV and wind must fall sharply very quickly; to produce this effect we just use very low initial low growth rates, much lower than observed historically.

For a given component, let the annual growth rate of production be g_t , where

$$g_t = \begin{cases} g_{t_0} & \text{if } 0 \leq t < t_1, \\ g_{t_0} + \frac{g_T - g_{t_0}}{1 + e^{-s(t-t_1 - (t_2-t_1)/2)}} & \text{if } t_1 \leq t < t_2, \\ g_T & \text{if } t_2 \leq t \leq T \end{cases} \quad \text{with steepness parameter } s = 50 \left| \frac{g_T - g_{t_0}}{t_2 - t_1} \right|. \quad (1)$$

Each growth-rate-determined component has its own set of growth rate parameters $\{t_1, t_2, g_{t_0}, g_T\}$ in each scenario. The start time is t_0 . There is a constant initial growth rate, g_{t_0} , which persists until time $t_0 + t_1 - 1$. From $t_0 + t_1$ to $t_0 + t_2 - 1$ the growth rate follows a logistic growth curve with steepness parameter s , reaching the constant final growth rate g_T at time $t_0 + t_2$, which persists until time horizon $t_0 + T$.[§] This is shown in Panel A of Figure S2.

Let annual production be q_t . Given growth rate trajectory g_t , the basic rule for calculating annual production is

$$q_{t+1} = (1 + g_t)q_t, \quad t = 0, 1, 2, \dots, T, \quad g_t > -1. \quad (2)$$

However, there are several exogenous constraints in the model, and for some variables it is necessary to impose an additional cap on production in order to meet a constraint. Let this cap be q_t^{max} . In order to control the behaviour of the cap, we introduce a new variable t_3 . This defines the time after which the cap becomes active. There are three cases (shown in Figure S2):

- (i) $t_3 \geq T$, so that the cap is never active. In this case production is determined purely by the growth curve in Panel A. This production is shown in Panel B.
- (ii) $t_3 = 1$, so that the cap is always active. In this case the cap begins and remains at or above production for all time. This is shown in Panel D – the production implied by A eventually hits the cap. The realized growth rate curve corresponding to the application of the cap is shown in Panel C.
- (iii) $1 < t_3 < T$, so the cap is only active after t_3 periods. This case is required to deal with a situation where there is an initial transient period in which fast production growth climbs above the cap, but then returns below it later on, so that in the long-term the cap still functions as an upper bound (Panels E and F).

The equation for production is therefore

$$q_{t+1} = \begin{cases} (1 + g_t)q_t & \forall t, \text{ if } t_3 \geq T \text{ (i.e. (i))} \\ \min(q_t^{max}, (1 + g_t)q_t) & \forall t, \text{ if } t_3 = 1 \text{ (i.e. (ii))} \\ (1 + g_t)q_t & \text{for } t \leq t_3, \text{ if } 1 < t_3 < T \text{ (i.e. (iii))} \\ \min(q_t^{max}, (1 + g_t)q_t) & \text{for } t > t_3, \text{ if } 1 < t_3 < T \text{ (i.e. (iii))} \end{cases} \quad (3)$$

Figure S2 demonstrates the process, with $t_0 = 2020$, $t_1 = 10$, $t_2 = 40$, $t_3 = 10$, $T = 50$, $g_{t_0} = 0.24$, and $g_T = 0.02$.

We use this method because it is consistent with how technologies have developed historically⁴⁷, and because it makes describing how scenarios are constructed straightforward. However, it is important to note that for each growth rate trajectory specified in this way, there is a cluster of similar growth rate trajectories, all of which result in very similar scenarios. Many of these can be viewed as functionally equivalent, extending the validity of our results beyond the five scenarios used in this analysis. For example, see Figure S78, which shows an alternative to the hydrogen growth trajectory used in the Fast Transition scenario. In our model, following this alternative trajectory would only increase the probability of electrolyzer costs decreasing faster, yielding a lower cost scenario, and hence these trajectories may be considered equivalent in the context of this study, as the main results would still be valid.

In fact, since the model is based on stochastic experience curves, the general heuristic is that our results are likely to remain valid for all alternative scenarios in which any high-learning technology is modified with *on average* higher, earlier growth. This includes alternative growth rate trajectories with potentially large fluctuations (as have been observed historically for solar, wind, batteries and electrolyzers, due to policy and other price/demand dynamics – see Panel D of Figure S8). Thus both the level and duration of the growth rates we use to construct scenarios (again see Panel D of Figure S8, plus the left-hand column of Figure S9) should not be viewed as hard-and-fast rules, rather as representative trajectories, representing whole clusters of scenarios with very similar characteristics.

[§]For simplicity, as shown in Eq. 1, a single steepness parameter coefficient of 50 is used for all components. This is not a particularly important modelling choice, and the specific value has little impact on the results.

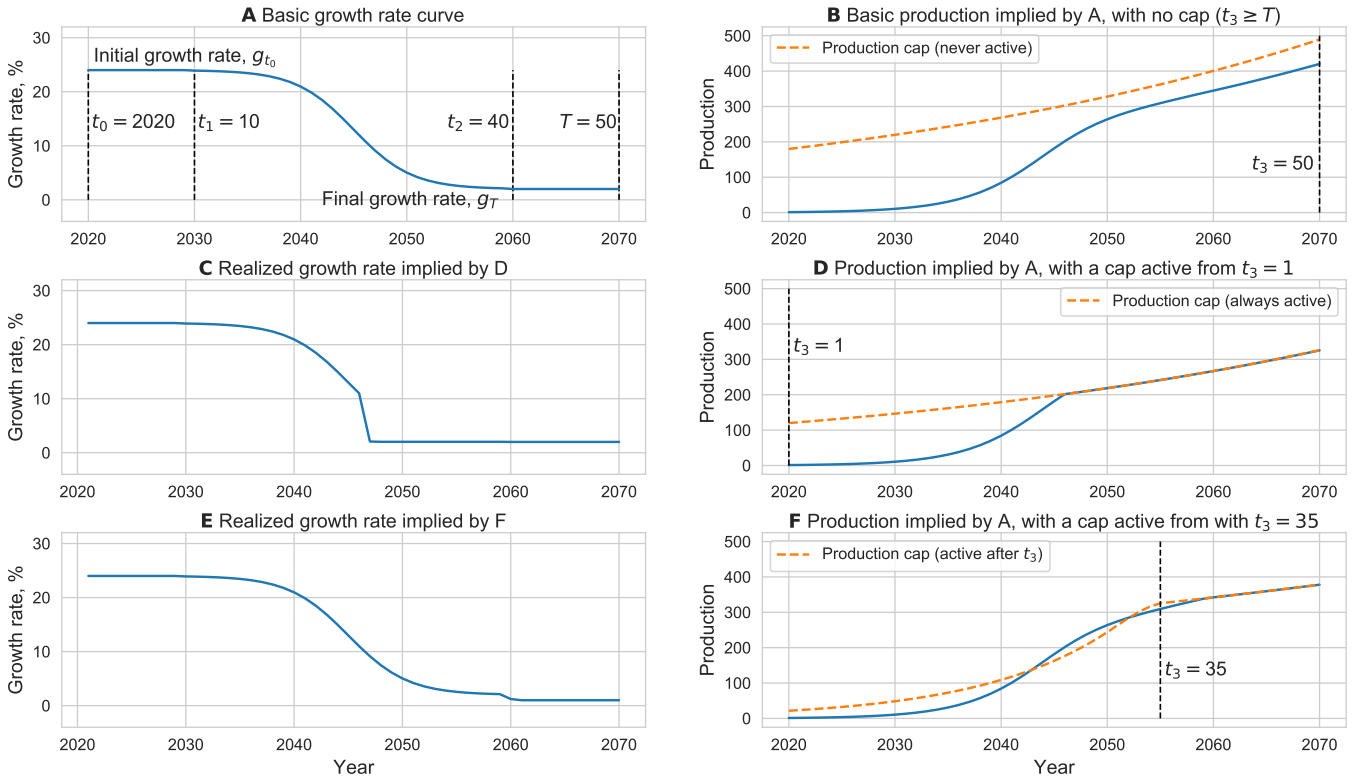


Figure S2: The growth process used in scenario construction. Panel A shows the basic growth rate curve, g_t , as defined by Eq. 1, which is used to construct a production curve, q_t . Possible production curves are shown in panels B, D and F, but which production curve is generated depends on whether there is a production constraint in place to cap growth, as defined by Eq. 3. Panel A shows the growth rate transitioning from 24% to 2% per year. As shown in Panel B, this corresponds to a huge expansion in production over 20 years, followed by a gradual slowdown to a new long-term steady state (this is case (i), i.e. Eq. 3 top row). Panel D shows production in a case where the growth eventually becomes capped by an exogenous constraint (case (ii), Eq. 3 second row). Panel C shows the implied actual growth rate of production in such a case; instead of following the smooth curve shown in Panel A the curve gets clipped as the production cap is reached. Panel F shows production in a case where a cap exhibits transient behaviour and only becomes active after time $t_3 = 35$ (case (iii), Eq. 3 third and fourth rows). Panel E shows the implied actual growth rate of production in this case. (Note that growth rate curves in panels C and E are not used to generate production curves, but rather are the consequence of combining the curve in panel A with a production cap.)

2.2 Energy services

Energy services have been the foundation of human prosperity, so we build the model around these. Our core assumption is that the provision of energy services increases at 2% per year until 2100. The IEA estimates that between 1973 and 2017 primary energy consumption grew at 1.9% per year, final energy consumption grew at 1.7% per year and electricity generation grew at 3.3% per year on average²⁴. Other sources estimate that primary energy consumption grew at 1.1% per year between 1800 and 1950 and 2.6% per year between 1950 and 2017 on average (or 1.5% over the period 1800 to 2017)⁴⁸. As mentioned earlier, predicting changes in conversion efficiencies between useful energy and energy services is beyond the scope here, so we just take useful energy as a proxy for energy services, and assume that useful energy grows at 2% per year.[‡]

We specify that the quantity of useful energy provided by each end-use sector is exactly the same in all scenarios, so that identical levels of energy services are supplied. This assumption also applies to the electricity component of the energy sector, but not to the fossil fuel components (because these depend on end-use fossil fuel consumption, so they must be scenario-dependent). The situation is shown in Figure S3. In contrast to many other energy models, in particular partial equilibrium models, in which energy service consumption varies depending on energy prices, our method allows for a simple, transparent and meaningful (“apples-to-apples”) comparison of the costs associated with different scenarios.

[‡]This leaves room for further reductions in useful energy consumption (and energy costs), since energy services may be provided more efficiently in future.

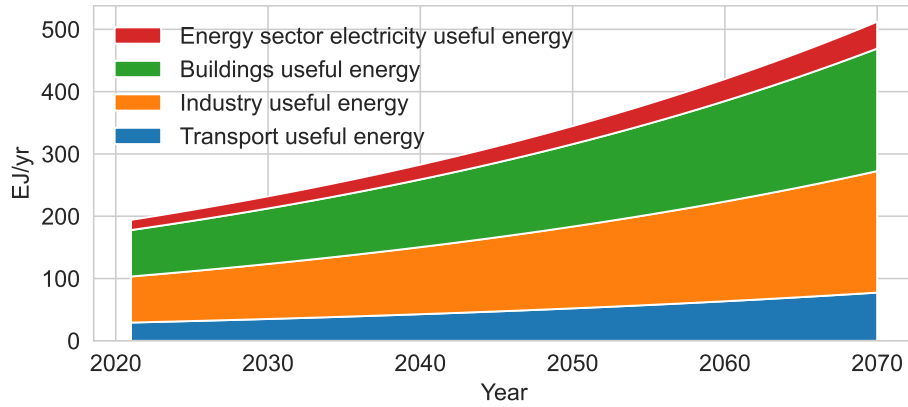


Figure S3: Exogenous assumption of useful energy demand growth.

2.3 Indices, slack variables and meeting the useful energy constraint

With the total useful energy supplied to each end-use sector completely determined in this way (plus energy sector electricity), we first grow and contract the different energy carriers to meet the useful energy totals in each year, then grow and contract the energy inputs (direct-use primary energy fossil fuels and electricity generation resources) such that the carrier totals are also met (given fixed conversion efficiencies).

The process can be conceptualised by working from right to left on Figure S1: useful energy of end-use sectors (plus energy sector electricity) grows at specified rates, then the various energy carriers are grown or contracted to meet the sector demand constraints, then, in turn, various technologies are grown and contracted to meet the carrier demand constraints.

In each case, in order to meet a specific demand constraint, all input components *except one* are grown or contracted using the method in S.I. 2.1. The “left-over” component is called the slack variable. Different slack variables are used in meeting different constraints, depending on the scenario, and are selected in advance during the scenario design process. They are used to supply any residual demand left unfulfilled by the sum of all other components (for more detail see Eq. 7 below).

Once the end-use sectors (plus energy sector electricity) have been supplied with sufficient energy carriers, and these energy carriers have been supplied by direct-use primary energy and electricity generation resources, a few extra steps must be taken to complete the scenario. First, energy sector fossil fuel energy carrier consumption must be calculated (as a function of fossil fuels in end-use sectors) and added to the scenario. Second, the quantities of facilitating technologies must be calculated and added to the scenarios (i.e. utility-scale grid batteries, batteries for the transport sector, electrolyzers, and electricity networks).

Any mix of direct-use primary energy sources, electricity generation technologies, and other facilitating technologies that meets the carrier constraints and useful energy constraints in each year is a valid scenario. The costs of such valid scenarios may be compared fairly with one another, as they provide equivalent levels of energy services. A “scenario” then is just a precise production schedule for every technology in the model; or put another way, it is a technology portfolio, made up of the quantity of each technology produced each year of the model. For N technologies and T periods each technology portfolio has dimension $N \times T$, each component of which is a closed interval of real numbers. Here $N = 14$ and $T = 80$ (since scenarios are constructed to 2100), so each portfolio has dimension 1120. This space is far too large to explore even reasonably thoroughly by computation methods, so the only possibility is to use heuristic methods to construct and compare a relatively small number of portfolios. We construct a scenario, or portfolio, as follows.

Let t be the period index and $t_0 = 2020$ be the start year. Let i be the technology index and q_t^i be the quantity of technology i produced in period t . Let j be the energy carrier index, and k be the sector index. Let \mathcal{J}_k be the set of indices of the carriers used in sector k . Let \mathcal{I}_j be the set of indices of the technologies used as input to carrier j . Values are shown in Table S7. (Note that in this construction process the energy sector consists only of energy sector electricity; energy sector fossil fuel energy carriers are dealt with separately afterwards, in S.I. 2.6)

Index	Category	Component	Input component index set
i	1	Technology	Oil (direct use)
	2	Technology	Coal (direct use)
	3	Technology	Gas (direct use)
	4	Technology	Coal electricity
	5	Technology	Gas electricity
	6	Technology	Nuclear electricity
	7	Technology	Hydroelectricity
	8	Technology	Biopower electricity
	9	Technology	Wind electricity
	10	Technology	Solar PV electricity
	11	Technology	Daily batteries
	12	Technology	Multi-day storage
	13	Technology	Electrolyzers
	14	Technology	Electricity networks
j	1	Carrier	Oil
	2	Carrier	Coal
	3	Carrier	Gas
	4	Carrier	Electricity
	5	Carrier	P2X fuels
k	1	Sector	Transport
	2	Sector	Industry
	3	Sector	Buildings
	4	Sector	Energy sector electricity

Table S7: Model components, indices, and sets of indices of carrier inputs to sector k (\mathcal{J}_k), and primary energy and electricity generation inputs to carrier j (\mathcal{I}_j). See Figure S1 also.

2.4 Energy carrier quantities for sectors

Consider sector k . Let $EF_t^{j,k}$ be the final energy provided by carrier j to sector k at time t (with known initial values as shown in Table S4). Let $EU_t^{j,k}$ be the associated useful energy at time t . As described in S.I. 1.8, let the conversion efficiency of final to useful energy of carrier j when used in sector k be η_j^k . Then, by definition,

$$EU_t^{j,k} = \eta_j^k EF_t^{j,k}. \quad (4)$$

Let $EU_t^{total,k}$ be the total useful energy provided by sector k at time t . This is the sum of the useful energy provided by each sector k carrier:

$$EU_t^{total,k} = \sum_{j \in \mathcal{J}_k} EU_t^{j,k} \quad (5)$$

The RHS values are known at time t_0 , as shown in Table S4, and $EU_t^{total,k}$ is calculated for all times using the method in S.I. 2.2. It is thus defined exogenously, as shown in Figure S3.

Next pick a slack carrier variable for sector k , denote its index by $j_{slack}^k \in \mathcal{J}_k$. Then for each non-slack carrier, $j \in \mathcal{J}_k \setminus \{j_{slack}^k\}$, we specify two things: *i*) growth rate parameters $\{t_1, t_2, g_{t_0}, g_T\}$ that define the carrier's useful energy growth process (as in Eq. 1), and *ii*) an upper bound to the proportion of the sector's useful energy that the carrier may provide at any time.* Specifying these parameters allows us to implement a growth process with a cap. Let $\{t_1^{j,k}, t_2^{j,k}, g_{t_0}^{j,k}, g_T^{j,k}\}$ be the growth process parameters and $g_t^{j,k}$ be the resulting growth rate. Let $\xi^{j,k}$ be the upper bound fraction of the sector's total useful energy that carrier j may provide. Then, using Eq. 3, case *(ii)*, the useful energy provided by carrier j in sector k is given by:

$$EU_{t+1}^{j,k} = \min \left\{ (1 + g_t^{j,k}) EU_t^{j,k}, \xi^{j,k} EU_{t+1}^{total,k} \right\}, \quad (6)$$

and the final energy is calculated with Eq. 4.

*Such upper bounds are required because we assume some basic level of diversity among energy carriers is necessary – for example, we can't electrify everything in a given sector, there will always be a requirement for some other energy carriers.

Note that a particularly important set of parameters are the $\xi^{4,k}$, as these specify the upper limits to electrification of each sector (in terms of useful energy). These are important because electrification is one of the few available options for decarbonising the economy, due to the highly efficient delivery of energy services by electricity, and abundance of nearly zero carbon energy sources. Electrifying each sector as much as possible is necessary for transitioning to a sustainable energy economy, but we must also ensure that other carriers are available for processes requiring physical characteristics not provided by electricity. Several recent studies have focused on quantifying the technical potential for electrification in different sectors, and we use these as guidance for choosing our $\xi^{4,k}$ values.

Regarding transport, Mai et al.⁴⁹ considered on-road transport in the US, using technologies that are either market-ready or near-market-ready. By building a detailed, bottom-up, stock-taking model of energy infrastructure, they found that in 2050, in a high adoption scenario, the shares of EVs in the transport sector could be: 88% of light duty cars, 81% of light duty trucks, 94% of buses, 52% of medium-duty trucks and 37% of heavy-duty trucks. Bearing in mind that some of these shares could still be rising in 2050, and that light-duty cars and trucks make up around 90% of distance driven, it is plausible that essentially all light- and medium-duty transport could be electrified over the next few decades. In addition rail and short-distance sea transport are easily electrified (and to some extent short-distance air travel too). This just leaves some heavy-duty trucking, long-distance sea transport and long distance air transport (the first two of which may turn out to also be relatively easily electrified, e.g. for shipping see Kersey et al.⁵⁰). We therefore set $\xi^{4,1} = 0.8$, so that a maximum of 80% of useful energy in transport is delivered by electricity. This choice allows for the full electrification of light duty vehicles, but conservatively assumes that most long distance vehicles, air travel and shipping will still require liquid or gaseous fuels, accounting for the remaining 20% of useful energy.

Regarding industry, Deason et al.⁵¹ reviewed several recent studies and found that the overall technical potential for electrification is high. Non-electric applications in industry predominantly involve using fuel to generate process heat. Most of the reviewed studies estimated that, technically, all industrial heat applications can be entirely electrified. The report also notes that “existing estimates for industry electrification in the long term (e.g. 2050) are largely qualitative or at best semi-quantitative, and are largely insufficient to develop quantitative estimates for long term electrification potential, beyond the 100% technical potential quoted in many studies”. Madeddu et al.⁴³ found that 78% of useful energy in Europe’s industrial sector is electrifiable with technologies currently available and 99% is electrifiable with technologies now being developed. We use a conservative upper limit to electrification of industry, and set $\xi^{4,2} = 0.75$. This choice allows for a large residual fraction of industrial processes to still rely on liquid or gaseous fuel after maximum electrification, representing 25% of all useful energy consumed in the sector.

Regarding buildings, the Deason et al.⁵¹ report also finds that almost 100% of energy use in buildings can be electrified. This value is much less speculative than in the case of industry because all the required technologies already exist at commercial scale. We therefore set $\xi^{4,3} = 0.9$. This still allows for 10% of useful energy to be supplied by gaseous fuels after everything else has been electrified; this could represent some hard-to-convert residential heating for example, and some cooking applications.

In summary, based on our current understanding of the technical substitutability of technologies and processes, the upper limits to electrification of each sector that we use are (given as a percentage of total useful energy provided by each sector): $\xi^{4,1} = 0.8$, $\xi^{4,2} = 0.75$, $\xi^{4,3} = 0.9$.[†] Bear in mind that these figures refer to useful energy, and since electricity is the most efficient energy carrier the percentages in terms of final energy will be lower.

Although it is certainly possible to represent all the varied circumstances of energy consumption in end-use sectors in much greater detail than this, we use this strategy because it is a simple, transparent way of aggregating the many complex engineering challenges and constraints involved.

The method we have designed allows each non-slack carrier to be grown or contracted appropriately for the type of scenario being constructed. All parameters for each carrier must be chosen so that the total does not sum to more than the sector’s total useful energy, $EU_t^{total,k}$, and so that all transitions are suitably smooth (i.e. the upper bound does not cut in too sharply for any carrier). This approach allows the modeller to create the approximate growth dynamics desired quickly and easily, while still maintaining identical levels of useful energy across different scenarios.

The slack carrier variable is then used to make up the difference between the sector total useful energy

[†]We also set $\xi^{4,4} = 1$, though this is not used because energy sector electricity growth is exogenous and constant across scenarios.

and the sum of the useful energy provided by non-slack carriers:

$$EU_t^{j_{slack}^k} = EU_t^{total,k} - \sum_{j \in \mathcal{J}_k \setminus \{j_{slack}^k\}} EU_t^{j,k}. \quad (7)$$

The quantity of final energy delivered by each carrier to sector k is then given by Eq. 4. By applying this procedure for each end-use sector (plus energy sector electricity) we construct the annual quantities, in EJ, of final and useful energy of all energy carriers required to meet the useful energy constraints.

2.5 Electricity generation technologies

The electricity mix is specified in almost the same way as was used above to determine energy carrier use in sectors. Let the total quantity of electricity generated in year t be $Elec_t$. This is the sum of electricity for use in sectors and in P2X fuels production (P2X fuels are used in end-use sectors and for power grid backup, as described in S.I. 3.3.3 and S.I. 3.5). Transmission and distribution losses are small (around 5% in OECD countries⁵²) and falling, due to ever more advanced equipment; these are included in the model as energy sector electricity consumption. The total electricity generated each year is the sum of generation from the individual technologies:

$$Elec_t = \sum_{i \in \mathcal{I}_4} q_t^i \quad (8)$$

The initial values at time t_0 are all known, as shown in Table S8.

Source	Electricity generation		
	TWh	EJ	Year
Coal	9911	35.7	2019
Gas	6356	22.9	2019
Nuclear	2790	10.0	2019
Hydroelectric	4236	15.2	2019
Biopower	709	2.55	2020
Wind	1596	5.75	2020
Solar PV	833	3.0	2020

Table S8: Initial electricity generation data, from IEA WEO 2021. (2019 values are used in cases where 2020 values are temporarily low due to the COVID-19 pandemic.)

First, pick a slack electricity generation variable, denote its index by $i_{slack} \in \mathcal{I}_4$. Next, for each non-slack technology, $i \in \mathcal{I}_4 \setminus \{i_{slack}\}$, we use Eq. 3, case (iii), to transition from its current state to some upper limit of the electricity mix. To do this we specify three things: 1) growth rate parameters $\{t_1^i, t_2^i, g_{t_0}^i, g_T^i\}$ that define the technology's growth process, 2) an upper bound to the proportion of total electricity it may provide, let this fraction be ψ^i , and 3) t_3 , the time value before which the upper bound fraction ψ^i does not apply, and after which it does. This is required because sometimes a technology initially has a large share of generation, and we want it to grow much slower than other technologies (so that its share decreases), but still end up being constrained by the share of the mix defined by ψ^i in the long run.

With g_t^i the growth rate resulting from Eq. 1 as before, the annual production of non-slack technologies is then

$$q_{t+1}^i = \begin{cases} (1 + g_t^i)q_t^i & \text{if } 0 < t \leq t_3, \\ \min\{(1 + g_t^i)q_t^i, \psi^i Elec_{t+1}\} & \text{if } t > t_3, \end{cases} \quad (9)$$

and the slack technology is subsequently given by

$$q_{t+1}^{i_{slack}} = Elec_{t+1} - \sum_{i \in \mathcal{I}_4 \setminus \{i_{slack}\}} q_{t+1}^i. \quad (10)$$

2.6 Energy sector fossil fuels

The energy sector uses both electricity and fossil fuel energy carriers to produce refined fossil fuel products. For example gasoline and many other fuel oils are made by extracting crude oil and processing it in refineries.

The electricity consumption component ($EF_t^{4,4}$) was dealt with in the above sections S.I. 2.3-2.5, but the fossil fuel energy carrier components ($EF_t^{1,4}$, $EF_t^{2,4}$, and $EF_t^{3,4}$) were not. We calculate these here, and add them to the model.

The refined fossil fuel products manufactured by the energy sector are used as direct-use energy carriers both in end-use sectors and the energy sector itself, and also as inputs for electricity generation. To model the energy needs of manufacturing these products in future, we assume that the quantities of fossil fuel energy carriers required by the energy sector scale directly with the quantities of fossil fuel energy carriers used in end-use sectors, plus those used to make electricity. Furthermore, because we do not have data detailing how much of each fossil fuel input is required to produce each of the fossil fuel energy carriers, we assume that the current proportions of oil, coal and gas used by the energy sector stay fixed through time. The initial quantities of oil, coal and gas energy carriers used by the energy sector are $EF_0^{1,4}$, $EF_0^{2,4}$, and $EF_0^{3,4}$. In 2018 these values were: 13.9 EJ of oil, 17.1 EJ of coal, 16.5 EJ of gas (making a total of 47.5 EJ), so the fixed proportions we assume for these fuels in future are 29%, 36% and 35%, respectively.

To calculate the quantities of fossil fuels required as inputs for electricity generation, we need conversion efficiency factors for coal and gas electricity. Let these be ζ_i (with $i = 4, 5$), so that x EJ of fossil fuel results in $\zeta_i x$ EJ of electricity. In 2018, 102.6 EJ of coal generated 36.4 EJ of electricity, giving a conversion efficiency factor of 36%. Similarly, 54.5 EJ of gas generated 22.0 EJ of electricity, yielding efficiency factor 40%. These values are current global averages, and are lower than values seen in some modern plants, so as old plants are retired and the fleet is modernised, we expect higher efficiencies. Consistent with expectations of future efficiency factors^{53,54}, we use a coal power efficiency factor of $\zeta_4 = 0.40$, and a gas power efficiency factor of $\zeta_5 = 0.50$.

Then, the total quantity of each fossil fuel required as direct-use fuel for end-use sectors, plus as input for electricity generation is:

$$\text{Oil } (j = 1): \quad \sum_{k=1,2,3} EF_t^{1,k} \quad (11)$$

$$\text{Coal } (j = 2, i = 4): \quad \sum_{k=1,2,3} EF_t^{2,k} + \frac{q_t^4}{\zeta_4} \quad (12)$$

$$\text{Gas } (j = 3, i = 5): \quad \sum_{k=1,2,3} EF_t^{3,k} + \frac{q_t^5}{\zeta_5}, \quad (13)$$

and the total quantity of direct-use end-use sector fossil fuels plus total electricity generation fossil fuels at time t is the sum of these:

$$FossilFuels_t^{(\text{end-use} + \text{electricity})} = \sum_{k=1,2,3} EF_t^{1,k} + \sum_{k=1,2,3} EF_t^{2,k} + \frac{q_t^4}{\zeta_4} + \sum_{k=1,2,3} EF_t^{3,k} + \frac{q_t^5}{\zeta_5} \quad (14)$$

We therefore take the energy sector final energy provided by fossil fuel energy carrier j to be

$$EF_t^{j,4} = \frac{EF_0^{j,4}}{EF_0^{1,4} + EF_0^{2,4} + EF_0^{3,4}} FossilFuels_t^{(\text{end-use} + \text{electricity})} \quad j \in \{1, 2, 3\} \quad (15)$$

The results of applying this method can be seen in the ‘‘Energy Sector’’ rows of Figures S7, S9, S11, and S13.

2.7 Direct-use fossil fuels

The quantities of oil, coal, and gas energy carriers required in a given scenario are then simply the sum of the final energy used in each sector, as calculated above:

$$q_t^1 = \sum_{k=1,2,3,4} EF_t^{1,k} \quad (16)$$

$$q_t^2 = \sum_{k=1,2,3,4} EF_t^{2,k} \quad (17)$$

$$q_t^3 = \sum_{k=1,2,3,4} EF_t^{3,k} \quad (18)$$

We have now specified the energy quantity requirements for all sectors, energy carriers, and energy input technologies ($i = 1-10$). The model components that remain to be quantified are the energy storage, conversion and distribution technologies ($i = 11-14$), which we address next.

3 Energy storage and flexibility requirements

Because solar and wind power are intrinsically variable, as they become the dominant sources of energy we must be able to store and move energy easily, in order to constantly balance supply and demand, and provide both grid electricity and all other energy services at the different levels of reliability required by different users. There are many ways to solve this problem, many of which are near-optimal^{55,56}, and solutions depend on the many diverse characteristics of both electric power systems and the wider energy system.

The strategy we implement in our model revolves primarily around adding technologies, in appropriate quantities, that are able to store energy on either short, medium or long timescales. Short timescales are defined as being up to a day; medium timescales are around one to three days, and long timescales as more than three days, though this may be weeks or months.

A key background reference for understanding energy storage and flexibility options is Blanco and Faaij⁵⁷. This is a comprehensive review of over 60 different studies and reports from the last decade, with a focus on long term storage needs in high renewables energy systems. Some of the most important ways to cope with balancing supply and demand of electricity over different temporal and geographic scales are:

1. *Interconnection*, i.e. network expansion – having enough transmission capability to bring power from distant sources to satisfy local excess demand.
2. *Energy storage technologies* – storing energy so that it can be used when needed; common methods of energy storage include batteries, pumped-hydropower, capacitors, compressed-air, flywheels, gravity-based systems and storing chemical energy in fuels (e.g. hydrogen, ammonia, methane).
3. *Sector coupling* – connecting the power system to other sectors of the energy system (transport, industry, buildings), in particular via transport services, heating services, and economy-wide demand for low carbon fuels.
4. *Flexible generation* – holding dispatchable sources of power such as gas- or P2X-fired electricity in reserve, so they can be quickly deployed when needed (potentially batteries also, depending on the time scale).
5. *Demand side management* – scheduling demand for electricity so that it occurs when there is available supply.
6. *Optimizing the wind/solar ratio* – choosing a mix that best utilizes natural correlations between supply and demand.
7. *Overcapacity* – installing excess VRE capacity so that during periods when capacity factors are below average, generation still meets some given target level (as specified in the planning process).

Two further concepts that are useful in understanding how grid flexibility options relate to system reliability are related to the characterisation of electricity supply and demand. On the demand side, fluctuations are usefully characterised in terms of their time scales: seconds-to-minutes, daily, seasonal and annual. Different flexibility options are appropriate for demand variation on different time scales. On the supply side, electricity generation resources are usefully characterised in terms of being either *i*) fuel-saving VRE resources, i.e. resources that save on costs and emissions associated with fuel combustion, but whose output inherently varies on the scale of seconds to years (often highly predictably); *ii*) fast-burst balancing resources, available on the scale of seconds to hours; *iii*) firm resources, available in all seasons over days, weeks or longer⁵⁸.

All the above concepts overlap and intersect in complicated ways, which is why system planning – or even just broadly characterising what are likely to be the most desirable solutions – is so challenging⁵⁹. It is noteworthy, though, that the lenses through which energy systems research is conducted, and the kind of scenarios explored, have changed a lot through time. Only since around 2017 have the near-full electrification of transport and heat, and the role of long-duration energy storage (LDES) begun to be studied in sufficient detail. Further, it is important to acknowledge the role of technology forecasts in shaping the research landscape. It appears likely that energy modellers did not routinely consider such scenarios because they relied on far-future cost projections that were much too high. Only since observed solar, wind, battery and electrolyzer costs have continued to drop rapidly in the last few years (following their long-run trends) have far-future cost projections used in models been lowered. In particular, cost projections made by the IEA and others acted as frames of reference for projections implemented in IAMs, resulting in a failure to explore

the space of feasible energy scenarios well enough. The need to explore “extremes” more systematically is now recognised⁶⁰ – though what is considered extreme clearly varies through time.

The dependence of storage and flexibility scenario exploration on cost forecasts highlights the need for rigorous, statistically validated technology cost forecasts, such as those used in the study. It also highlights the fact that study results should always be interpreted in the context of the economic assumptions made. For example, in their 2018 study on firm low-carbon generation sources, Sepulveda et al.⁵⁸ use PV capital costs taken from (or in one case derived from) NREL’s 2017 Annual Technology Baseline (ATB): 1800 \$/kW (“Conservative”, based on then present costs), 900 \$/kW (“Mid-range”, also based on then present costs), and 670 \$/kW (“Very low”, based on far future costs). In a follow-up 2021 study on the relationship between firm generation and LDES, Sepulveda et al.⁶¹ use values from NREL’s 2018 ATB: 1011 \$/kW and 611 \$/kW (both based on 2045 fixed axis PV costs). At the time of writing (August 2022), NREL ATB shows 2050 PV capacity cost values of 752 \$/kW (“Conservative”), 618 \$/kW (“Moderate”), and 466 \$/kW (“Advanced”), and the global average cost of PV is 857\$/kW⁶². This wide variation in costs, depending on both the forecast year and forecast horizon, highlights the challenge modellers face in trying to select appropriate cost values to use in their models. In this case the “Mid-range” PV capital costs used in a 2021 study (i.e. 2045 cost projections from NREL’s 2018 ATB) are already above the current global average cost. If highly unrealistic future values are used in models, then this may reduce the relevance of their results.

As the above example shows, the literature is currently moving quickly, and many results and scenarios may be rendered obsolete within just a few years of publication. It is therefore important, in considering the storage and flexibility requirements of scenarios in our model, to try to pull out only those key results from the literature that are likely to be relatively robust to changes in cost assumptions, and not be too dogmatic about many other details. Related to this is the fact that our model is far too low resolution to be able to reflect all aspects of the storage and flexibility results found in the literature, so our strategy is to approximately map existing results to equivalent results applicable in our model.

The following key recent studies, which focus on specific aspects of the problem, were used in designing our model, and are relevant to support the adequacy of our energy storage and flexibility assumptions:

1. *Sector coupling* – Brown et al.³⁷, Victoria et al.⁶³, Gea-Bermúdez et al.⁶⁴, Kittner et al.⁶⁵.
2. *Role of long duration energy storage* – Ziegler et al.⁶⁶, Dowling et al.⁶⁷, Albertus et al.⁶⁸, Guerra et al.⁶⁹, Hunter et al.⁷⁰.
3. *Firm low carbon generation* – Sepulveda et al.⁵⁸, Sepulveda et al.⁶¹.
4. *Seasonal storage* – Victoria et al.⁶³.
5. *Geophysical constraints on the reliability of VRE sources, and the relation to storage, overcapacity and inter-connection* – Shaner et al.⁷¹, Tong et al.⁷².
6. *Geographical and temporal variation in VRE sources* – Grams et al.⁷³, Zeyringer et al.⁷⁴, Jerez et al.⁷⁵, Tröndle et al.⁷⁶.

As summarised by Gea-Bermúdez et al.⁶⁴, once very high VRE penetration and full sector coupling are included in a model, this enables a paradigm shift in supply-demand dynamics. The system then shifts from the current state, in which electricity generation adapts to meet changing demand, to one in which an expanded “demand” concept – including demand for both direct energy services and energy to recharge large stores embedded within different sectors – adapts to natural variations in generation. This is the basis for our Fast Transition scenario. We now discuss the specific storage and flexibility choices we have made in our model.

We do not model specific demand-side management technologies (such as flexible EV charging and residential smart meters) explicitly because very little data exists and their progress is hard to model (i.e. there is no good cost or deployment data that allows us to apply Wright’s law (see S.I. 5.2)). However, demand-side management is already widely used in industry and very large buildings, and is becoming increasingly attractive in residential settings, with operators able to provide virtual power plant services to the grid. It is a relatively cheap way of matching supply and demand⁷⁷. We expect that demand-side management technologies will exist on a vast scale as devices become ever more connected and responsive, particularly as electric vehicles become widespread. Our transition scenarios are therefore pessimistic, in the sense that they provide more (relatively expensive) storage than would be necessary if demand-side management were included in the model.

Interconnection can be a good substitute for storage or backup generation during periods (hours to weeks) when VRE capacity factors are locally low due to local variation in weather patterns. This can allow power to be transported in from distant locations with different weather and higher capacity factors. Studies often find it to be a cheap solution to cope with the intermittency of renewables, e.g. Brouwer et al.⁷⁷. In addition, interconnection may often be required to transfer bulk power from large, remote VRE generators over long distances to load centres (and there are usually additional benefits in terms of grid resilience). Unfortunately, like demand-side management, it is hard to quantify cost savings from interconnection in terms of technology deployment costs, so that we cannot model its likely progress using Wright's law. We do include grid investment costs in our model, as described in S.I. 3.7, and some of these correspond to expanding the transmission system, so they may be considered as interconnection investments. However, we do not model grids or interconnection in detail, as this is incredibly complex. We instead focus on modelling bulk energy storage, in the knowledge that improved interconnection may be a much cheaper solution in many cases. Highly complex models will be needed to manage the design and expansion of power grids, carefully taking in to account the technical characteristics of individual grids based on the geography of both their loads and renewable resources (see e.g. Zeyringer et al.⁷⁴, Denholm and Mai⁷⁸, Solomon et al.⁷⁹, Saarinen et al.⁸⁰, Ziegler et al.⁸¹, Victoria et al.⁶³).

Current solutions to matching fluctuating supply and demand in electricity markets depend strongly on maintaining a reserve capacity of dispatchable generators. Since all of our scenarios begin with the present mix, they all make use of this at the outset. However, our Fast Transition scenario phases out coal and gas powered electricity fairly quickly, and after 2050 it has only minimal presence. Instead dispatchable power is based on stored P2X fuels that are produced from renewables, plus relatively small amounts of hydropower, nuclear, and biopower.

The distinction between reserve capacity and energy storage technologies becomes blurred once P2X fuels are considered, as P2X fuels are used to store energy, and are used together with turbines or fuel cells to dispatch electricity when needed. Historically, the ability of fossil fuel generators to ramp up and down as required by the system operator was the main form of very short term grid flexibility; we assume batteries will perform this service. More recently, as solar and wind have grown, long periods of reduced output from these sources (potentially up to a few weeks) have also been accommodated by simply ensuring adequate stocks of fossil fuel. We use P2X fuels to replace these.

Given the tremendous complexity of options available for dealing with intermittency, and the large number of degrees of freedom, there is no unique method for ensuring adequate levels of technology are present to provide society with highly reliably electricity. The method we use is as follows. It is known that for large deployments of solar, a large capacity of daily storage is essential for overcoming the daily solar cycle⁷². This is less critical for wind due to its geographical diversity, but small amounts of daily storage capacity still yield large benefits. We consider the evidence for how much daily storage capacity is required and include this in the model, as a function of the percentage of VRE in the generation mix (*nb.* this is not the same as the percentage of VRE in the final electricity mix). In addition, long periods of reduced VRE output are likely due to natural weather patterns, so we want to ensure that the longest likely such period is covered by some form of alternative generation source. This may be considered seasonal storage, as it involves shifting large quantities of energy over periods of weeks or more. We consider the evidence and ensure this capability is present as necessary. Finally, increased diversity of flexibility options is known to increase security of supply and lower costs⁵⁷, so we include low levels of one other technology option, multi-day storage. Once we have determined the quantities of technologies required to fulfil these requirements, we verify that the final combination is likely to provide very high levels of reliability by comparing the resulting system properties with the criteria set out by Tong et al.⁷² in their study of the geophysical availability of solar and wind resources.

Our Fast Transition scenario thus makes extensive use of energy storage on three different timescales. As described in more detail below, as representative technologies, we choose lithium-ion batteries for short timescales, redox flow batteries for intermediate timescales and P2X fuels for long timescales. The technologies we use here have the advantage that we have time series of cost and production that we can use to apply Wright's law. We do not claim that these are the cheapest solutions – we doubt that they are, at least taken alone – but they provide an upper bound on the likely costs. In fact, it should be noted that our modelling strategy is chosen largely to take advantage of the available data and for the simplicity of model implementation and communication. It is even possible that bulk energy storage, as envisioned here, could be one of the costliest methods of managing VRE sources⁴⁵. Thus our approach should be viewed as sufficient, comprehensible, and providing a conservative upper bound on the cost of integrating large shares of VRE sources in the energy system.

3.1 Short-term batteries

We assume that daily cycling batteries (which we also call short term batteries) are used both to manage daily fluctuations in the supply and demand of electricity and to provide energy storage for electric vehicles. We assume that Li-ion batteries are used. These have a strong record of progress historically⁸², and there is no reason to think this shouldn't continue. However, it is quite possible that better solutions may be found in the future.

Let the quantity of short-term batteries ($i = 11$) used in power grid applications be $q_t^{11,grid}$, and in transport be $q_t^{11,transport}$, so that

$$q_t^{11} = q_t^{11,grid} + q_t^{11,transport}. \quad (19)$$

First consider power grid batteries. The only nondispatchable electricity generation sources in the model are wind ($i = 9$) and solar ($i = 10$). The total quantity of electricity generated by these technologies is $(q_t^9 + q_t^{10})$ EJ each year, or $\frac{1}{365}(q_t^9 + q_t^{10})$ EJ each day, on average. Storage requirements are defined in terms of this average daily VRE generation. We specify that enough short-term battery capacity must be available to store a fixed fraction β_{ST} of this energy. We set $\beta_{ST} = 0.2$ in high renewables scenarios and 0 otherwise.

Note that β_{ST} is the fraction of electricity *generated* by wind and solar, not the fraction of final electricity *delivered* by wind and solar via the grid. In a high renewables scenario VRE generation will be much higher than VRE delivered directly via the grid, because so much electricity is required for P2X fuel production in the energy sector (perhaps 50% higher, see Panel H of Figure S8).

To illustrate the point, suppose that total generation is exactly 50% higher than final electricity, variable renewable energy constitutes 85% of total generation, and $\beta_{ST} = 0.2$. In this case short-term battery capacity will be available to store 25.5% of average daily final energy (i.e. $0.2 \times 0.85 \times 1.5$), corresponding to load-shifting 6.12 hours of final electricity each day. This is longer than the typical evening peak in electricity demand.

Solomon et al.⁷⁹ used simulated wind and solar generation data, with hourly electricity demand data, to model the Californian power grid in a very high renewables scenario. They found that network storage corresponding to around 22% of the average daily demand was sufficient to meet the hourly demand year round, and increasing the storage above this showed decreasing benefits. VRE penetrations of up to 85% were achieved with this level of storage when energy loss was allowed (i.e. due to excess generating capacity). Our value of $\beta_{ST} = 0.2$ exceeds this level of storage.

Now, the current installed capacity of short-term grid batteries is very small, so trying to impose the β_{ST} condition immediately would require a many-fold increase in capacity instantaneously in high renewables scenarios. This is unrealistic, so instead we grow the installed storage capacity at a constant rate (approximately equal to the current growth rate) until the β_{ST} condition is first met, and from then on use β_{ST} to track the variable renewable energy penetration as described above. Let the quantity of short-term batteries initially grow with constant annual growth rate g_{ST}^{grid} . Then

$$q_{t+1}^{11,grid} = \min \left\{ (1 + g_{ST}^{grid}) q_t^{11,grid}, \frac{\beta_{ST}}{365} (q_{t+1}^9 + q_{t+1}^{10}) \right\}. \quad (20)$$

Both the growth rate, denoted by g_{ST}^{grid} , and the fraction β_{ST} are scenario-specific.

Next consider transport batteries. The total quantity of electricity delivered to the transport sector each year is $EF_t^{4,1}$ EJ. This equates to $\frac{1}{365} EF_t^{4,1}$ EJ per day on average. We assume that at the start of each day all EV batteries are full and at the end of the day they are empty. This crude assumption accounts for neither the possibility of charging at multiple locations throughout the day, nor the complex relationship between charging dynamics, battery degradation and longevity, and the ability to offer vehicle-to-grid (V2G) services. We neglect all these factors. Complicated battery management systems will be required. The ability to charge at multiple locations throughout the day would reduce the size of battery pack required for a given level of daily usage, but we do not consider this.

Our assumption of one full charge/discharge cycle per day implies that there must be $\frac{1}{365} EF_t^{4,1}$ EJ of EV battery pack capacity on the system at time t . As with grid batteries above though, imposing this condition immediately would result in a large instantaneous jump in installed capacity, so again we assume the current growth rate continues until the condition is met. Let the initial constant annual growth rate of these batteries be $g_{ST}^{transport}$. Then

$$q_{t+1}^{11,transport} = \min \left\{ (1 + g_{ST}^{transport}) q_t^{11,transport}, \frac{1}{365} EF_{t+1}^{4,1} \right\}, \quad (21)$$

and q_t^{11} is now fully determined.

The methods we use here take account of the round-trip efficiency of batteries. First, in the case of transport batteries, this is because the final to useful energy conversion efficiency of electricity in the transport sector is 0.8 (see Table S4), so electricity requirements – and hence battery capacity too – have already been scaled up to account for efficiency losses. Second, in the case of grid batteries, it is because the cost metric we use to account for the costs of deployed battery systems (\$/kWh) already accounts for battery efficiency. As described in Cole et al.⁸³, battery system costs are given in terms of dollars per usable kWh of battery storage, and “usable kWh of battery storage means that round trip efficiency and depth of discharge are accounted for in the price of the battery pack in dollars per kWh”. Therefore no further modifications are required.

3.2 Multi-day storage

We use flow batteries as a representative concept for multi-day storage. These have modular, extensible energy storage reservoirs, which make them a good candidate for bulk, multi-day (or even multi-week) storage, with potentially good economies of scale* (see e.g. Ha and Gallagher⁸⁴, Li et al.⁸⁵). In order to be specific and use data relating to an existing technology, we consider Vanadium redox flow (VRF) batteries in particular. These are at early commercialization stage, so there is only limited cumulative capacity and cost data that we can use to calibrate Wright’s law.

As with short-term batteries, we assume that enough multi-day storage capacity is available to store some fixed fraction, β_{LT} , of the average daily variable renewable energy generation, $\frac{1}{365}(q_t^9 + q_t^{10})$ EJ. On days with very productive weather, variable renewable energy generation covers demand and excess energy can be stored and released over the next few days. We set $\beta_{LT} = 0.1$ for high renewables scenarios and 0 otherwise. Then the installed capacity of multi-day storage ($i = 12$) in year t is

$$q_{t+1}^{12} = \min \left\{ (1 + g_{LT}^{grid})q_t^{12}, \frac{\beta_{LT}}{365}(q_{t+1}^9 + q_{t+1}^{10}) \right\}. \quad (22)$$

In the example above (S.I. 3.1), in which total generation is exactly 50% larger than final electricity and variable renewable energy constitutes 85% of total generation, a value of $\beta_{LT} = 0.1$ would correspond to the ability to load-shift 3.06 hours of energy over several days (or indeed every day if required). This is not a large amount of storage, but we include it to increase the diversity of options available. We assume that flow battery system costs are given in terms of dollars per usable kWh of battery storage (as is the case for Li-ion batteries above), so that battery system efficiencies are automatically included in the model.

3.3 P2X fuels for long term energy storage and hard-to-electrify applications

We use *Power-to-X* (P2X) fuels for long term energy storage and to serve energy applications that are hard to electrify. Here we take this to mean the conversion of electricity from wind and solar PV to any of several possible fuels, including hydrogen, ammonia, methane, methanol or others. (If carbon-containing fuels are used we assume this can be done so that it is carbon-neutral). This is typically done by using electrolyzers to convert electricity and water into hydrogen, and then if needed using the hydrogen that is produced to make other fuels that are denser and easier to handle. These processes all exist commercially today, so there is no doubt that they work and can be scaled up as necessary (see S.I. 8.7.1 for a discussion on the viability of the scale up of electrolyzers required in the Fast Transition). The interesting development expected now though is the introduction of cheap, scaleable renewable electricity to the process, which is likely to spur further research and innovation in the field.

We assume that P2X fuel is used for two purposes. The first is for direct use in end-use sectors, to provide a substitute for fuels such as diesel fuel in transportation, or coal in industrial applications where high-temperature heat is needed, such as steel manufacture. The second purpose is to produce electricity when there is a mismatch between supply and demand that can not be solved by battery storage, either due to long spells of reduced VRE output, or seasonal variation in demand patterns, or both.

There are currently two leading technologies available for converting hydrogen to electricity: fuel cells and direct combustion in gas turbines. Regarding the first of these options, one possibility is that reversible electrolyzer/fuel-cells become widely available, so that the same fuel cell both makes electricity from hydrogen and converts hydrogen back into electricity when needed. Of the three main electrolyzer technology

*However, their dependence on physical pipes and fluids may be a drawback when it comes to manufacture, installation and maintenance, relative to more self-contained electronic technologies such as capacitors or other battery chemistries. This remains to be seen.

options currently available (AEL, PEM, SOEC, see S.I. 6.13), SOEC is reversible; alkaline and PEM fuel cells do exist but as separate products from electrolyzers. SOEC is currently the most expensive of the three technologies, but its cumulative experience is very small (as is that of PEM), and its cost may well drop in line with the others (or perhaps innovation will allow the others to develop reversible functionality). This bi-directional setup is plausible because in a high renewables scenario a large quantity of P2X fuels will be produced when excess renewable energy is available, but when backup is required the opposite will be true, so that the fuel cell functionality will be available. Note that the quantity of P2X fuels for end-use sectors would be much larger than the quantity of P2X fuels for grid backup in this case, so the requirement for electrolyzer capacity would be much higher than for fuel-cell capacity. Since this eventuality is quite plausible, we could assume that enough of the electrolyzers already included in our model are reversible (about one quarter would be sufficient, see Figure S8 Panel H), so that there would be no extra cost involved in this method of electricity generation. However, we do not rely on this, as there is another solution.

The other possibility is that gas combustion turbines are available for generating electricity from hydrogen. Direct combustion in gas turbines is appealing because natural gas power stations can easily be converted to burn hydrogen instead, and there is very large capacity available. Gas turbine manufacturers already sell turbines with partial hydrogen co-firing capability, and several major suppliers are in the process of transitioning to 100% hydrogen burning capability.[†] Ammonia can also be combusted in turbines, though combustion must be complete in order to avoid undesirable combustion products⁸⁷.

We assume that gas turbines are used to generate electricity from P2X in our model. There is currently a very large fleet of gas turbines, which will be required for at least the next 15 years in all scenarios. During this period, as capacity is maintained or upgraded, it will become P2X-ready, so that as the transition proceeds, a very large quantity of generation capacity will remain available. Furthermore, over this period, all capital investments will be amortized, so that after around 2035, no capex costs remain. Recall, fuel costs, which are the major cost component of gas-fired electricity currently, and are likely to remain so in future, are already covered in the model. Thus we make the simplifying assumption that there are no further capex or opex costs that we need to represent to capture this energy conversion.[‡]

Perhaps in reality a combination of reversible fuel cells and combustion turbines will be used, depending on location and specific system requirements. Whichever path ends up being taken, we can assume that all associated capital costs are already included in the model, either in the electrolyzer costs or the fully amortized gas turbine costs. It would be possible to extend the model further to include some relevant cost component here explicitly, but we exclude it now for simplicity. We believe these costs are so small relative to other model components that they would barely impact the final results.

Regarding storage, transmission and distribution of P2X fuels, many options are in use already today, and we argue that total infrastructure costs would be no higher than total fossil fuel infrastructure costs in an equivalent scenario. Hydrogen is cheaper to make than denser P2X fuels, but more expensive to handle, so whether or not to upgrade it to ammonia or methane etc. will depend on the exact application and circumstances. Hydrogen liquifies at -253°C , so liquifaction is usually not an economical route, but ammonia liquifies at -33°C , which is very manageable. We summarise the most important points here, but for many more details see IEA²⁶, IEA⁸⁹, WEC²⁷, Blanco and Faaij⁵⁷.

3.3.1 P2X fuel storage

Hydrogen is currently stored as a gas or liquid in tanks for small-scale applications, but geological storage, which is used for natural gas today, would be required in a high hydrogen scenario. Salt caverns are the best solution for bulk hydrogen storage and have been used to store gaseous hydrogen since the 1970s, at high efficiency (98%) and very low cost (around 0.2 \$/kWh, at least an order of magnitude less than pumped-hydropower storage⁵⁷). An average size salt cavern is able to store 3 TWh of hydrogen, and Germany alone has 170 salt caverns that are currently used to store natural gas. The global potential for salt cavern capacity is well above global energy storage requirements. However, they are not available in all locations, so alternatives would also be required. Depleted fossil fuel reservoirs and aquifers are already used extensively for natural

[†]All EU turbine manufactures have committed to 100% hydrogen capability by 2030⁸⁶.

[‡]As a brief reality check, observe that the EIA's 2021 AEO⁸⁸ reports that for combined cycle gas turbines in the US, out of a total LCOE of 34.51\$/MWh, levelized capex accounts for 7.00\$/MWh (20.3%), levelized fuel costs are 24.97\$/MWh (72.4%), levelized transmission costs are 0.93\$/MWh (2.7%), and levelized fixed operations and maintenance (O&M) costs are 1.61\$/MWh (4.7%). The latter is the only component that would not be included in our model after a transition to P2X fuels. In the Fast Transition, approximately 5 PWh of final electricity comes from P2X fuel per year after 2050 (and almost none before that, see Figure S8 Panel K). At 1.61\$/MWh for levelized O&M costs, this would translate to around \$8 billion per year, which is around 0.1% of the expected \$6-8 trillion of annual expenditures included in the model (see Figure S45).

gas storage, and this may also be possible for hydrogen storage, but this is uncertain due to lack of experience with this method. If low cost geological hydrogen storage is not an option in a given region then ammonia, methane, or methanol may be good solutions.

Ammonia is much denser than hydrogen, and is generally much easier to handle and store. Typical large tanks store around 30,000 Mt, or 190 GWh (in terms of reformed hydrogen), and storage is very low cost, at around 0.1 \$/kWh⁹⁰, making it an excellent candidate for a large-scale energy carrier. The global ammonia supply chain already manages very large volumes and could be the basis for the huge volumes of P2X required in a zero carbon energy system.

3.3.2 P2X fuel transmission and distribution

Pipelines are efficient for bulk transportation at short to medium distances (up to 1500km), and there are thousands of kilometres of hydrogen and ammonia pipelines already in use around the world. This is an especially good solution for industrial hubs, where very large quantities of energy are used and technology coordination is easier. A well-established network of hydrogen pipelines already exists in North-West Europe, originating at the Port of Rotterdam, which is reshaping itself to become a low carbon industrial hub.

The total cost of transmission of hydrogen by pipeline is around double that of natural gas, per unit of energy⁹¹. This is due to the variable running costs of the extra compressors required to deal with the low density of hydrogen, and the capital costs of the compressors themselves. However, pipeline construction costs are about \$1million/km for both natural gas⁹² and hydrogen⁹³, as they depend mostly on labour, terrain and location. Construction costs vary greatly due to these three components alone⁹⁴. Modern natural gas distribution networks can also be used to transport hydrogen with only minor upgrades^{95,96}. A recent study found that retro-fitting the existing low pressure gas distribution network in Leeds, UK, would cost around £10,000 per kilometre of pipeline, far less than the cost of new pipelines⁹⁷. New ammonia pipelines are cheaper than new hydrogen pipelines.

For longer (intercontinental) distances, ammonia is already transported by sea in chemical tankers and liquid petroleum gas tankers. These are more expensive than basic oil tankers due to refrigeration requirements. Liquified hydrogen can be transported in ships similar to today's current LNG tankers (though presumably at similar, relatively high cost), and the first such ship was launched in Japan in 2019⁹⁸.

In summary, all of technologies required for storing and distributing hydrogen, ammonia and other P2X fuels already exist, mostly in commercial operation, though at much smaller scale than the existing fossil fuel network. There may be cost penalties associated with certain modes of storage and transport due to hydrogen's low density (e.g. long term storage in small quantities), so patterns of use and storage may change, though even this is uncertain when one considers the potential for technological improvements over the next few decades (e.g. Jensen et al.⁹⁹ and Butera et al.¹⁰⁰ consider reversible SOEC electrolysis with integrated methanation and natural gas storage). In any case, once scaled-up, it is likely that overall storage and distribution costs per unit of energy delivered would be broadly the same as those for fossil fuels delivered through the existing network, in both cases accounting for only a small fraction of delivered energy costs. Nevertheless, we conservatively assume that P2X infrastructure costs remain at their current levels, 2-3 times the cost of fossil fuel infrastructure, so that a P2X distribution network of a given size costs the same as an oil and gas distribution network 2-3 times the size.

The P2X sector in our Fast Transition scenario is, in 2070, just under half the size of the combined oil and gas sectors in the No Transition scenario, so it is safe to assume that total infrastructure installation, modification and upgrade costs for P2X fuels are *at most* roughly equal to total maintenance, upgrade and expansion costs incurred in the No Transition scenario. We therefore exclude both fossil fuel and P2X fuel infrastructure costs from the model.

3.3.3 P2X fuel quantity

We now calculate the quantity of P2X fuel required in each scenario. Recall that the maximum proportion of nondispatchable generation sources (wind and solar) in the electricity mix in the long run is given by $\psi^9 + \psi^{10}$ (these values are scenario-specific). We now define a new scenario-specific variable, θ , which is the fraction of final electricity generated annually from wind and solar that must be delivered via P2X fuels rather than directly from a generating source. This is because we expect there to be long periods of reduced solar and wind output that must be compensated for by some other means of generation. If in a given year the quantity of final electricity is X EJ, then a maximum of $(\psi^9 + \psi^{10}) \times X$ EJ of final electricity comes from wind plus solar. We specify that, after all major transition dynamics have occurred (i.e. from around 2060

onwards), $(1-\theta) \times (\psi^9 + \psi^{10}) \times X$ EJ of electricity comes directly from generation or batteries, but the remaining $\theta \times (\psi^9 + \psi^{10}) \times X$ EJ must be generated from P2X fuels *made from wind and solar*. We must then generate enough extra electricity from these sources to be converted to P2X fuels to cover the $\theta \times (\psi^9 + \psi^{10}) \times X$ EJ of final electricity.

The choice of θ depends on the scenario. Historical weather data for Great Britain shows that a lull in wind output of three weeks is around the maximum challenge a very high renewables grid there would likely face¹⁰¹. Grams et al.⁷³ classify different weather regimes in the Atlantic-European region, and consider the resulting large variations in wind and solar output, that last “several days or a few weeks” and affect neighbouring countries. However, they show that by locating European wind farms strategically far apart, effectively in different weather subsystems, the variability in total output can be almost entirely eliminated. Using such data to inform wind power planning decisions is an effective way to manage output on multi-day and longer timescales.

We assume that such results are broadly representative of all global regions. We assume that multi-week lulls in VRE output may occur, and so we specify that in a very high renewables scenario there must be enough P2X fuel stored to cover one month’s worth of final electricity each year ($\theta = 1/12$). In a medium renewables scenario, only one week must be covered ($\theta = 1/52$), and in a low renewables scenarios no P2X fuels coverage is required ($\theta = 0$).

Currently no P2X fuels are used in the grid, so we need to gradually phase in the quantity of final electricity provided via P2X fuels. For this we use our growth model again (Eqs. 1 and 3), with initial value 0, final value 1, and a phase-in period of $t_{P2X\ grid}$ years. This scenario-specific time value is chosen to roughly reflect the speed at which renewables are phased in. Let the resulting phase-in curve be called $\Pi_t^{P2X\ grid}$. The quantity of electricity that must be provided using P2X fuels as a source is therefore

$$\theta (\psi^9 + \psi^{10}) EF_t^{4,total} \min \left\{ \Pi_t^{P2X\ grid}, 1 \right\}. \quad (23)$$

Both methods that we consider for converting P2X fuels to electricity currently have conversion efficiency of around 50% (combusting hydrogen in gas turbines has the same efficiency as when burning natural gas, and see Staffell et al.⁹⁶ for fuel cell efficiencies), so the quantity of P2X fuels required is twice the quantity of electricity given in Eq. (23). We conservatively assume that this conversion efficiency does not improve, and thus assume that the total quantity of P2X fuels required throughout the system in year t is the sum of all end-use P2X fuels and grid backup P2X fuels:

$$P2X_t = \underbrace{EF_t^{5,total}}_{\text{Direct use}} + \underbrace{2 \times \theta (\psi^9 + \psi^{10}) EF_t^{4,total} \min \left\{ \Pi_t^{P2X\ grid}, 1 \right\}}_{\text{Electricity backup}}. \quad (24)$$

3.4 Verifying electricity system reliability

Shaner et al.⁷¹ performed a detailed study assessing the geophysical limits on solar and wind generation in the US, based on 36 years of weather data. A recent follow-up study using 39 years of resource data extended the analysis to 42 countries worldwide⁷², and we primarily use this to verify the reliability of our modelled system. They investigated how grid reliability depends on (1) the wind/solar generation mix, (2) annual wind/solar over-generation (above annual final electricity demand), (3) storage of varying durations, (4) interconnection (aggregation of resources over large areas), and (5) sensitivity of the results to changes in demand pattern in future, in particular a high electrification scenario. The study assumes perfect electricity transmission, with no losses or constraints, so, as with all modelling exercises, does not represent real-world conditions perfectly. Nevertheless, the results likely provide a very good indication of the approximate requirements, so we use them to verify that our assumptions regarding quantities of storage and electricity generation are likely sufficient to provide very high levels of grid reliability.

The study showed that without any storage or excess generation, solar and wind (in the appropriate ratio) can likely meet 72-91% of final electricity demand in all countries considered (83% on average). In this case the remaining 9-28% of demand must be met by dispatchable generation sources. However, adding 3 hours of storage (i.e. 12.5% of average daily electricity demand), allows solar and wind to meet 80% or more of demand in all but 1 country. Adding 12 hours of storage (50% of daily demand) pushes this figure to around 82% for all countries (though for most the value is around 90%). If instead (again with zero storage), one installs enough solar and wind capacity such that their annual electricity generation is 150% of the aggregate annual final electricity, then at least 82% of final energy can be met by solar and wind alone (with most

countries at 90% or more). A similar picture emerges in the case of adding interconnection between countries and regions. When these grid flexibility options are combined (in various combinations), solar and wind are easily able to meet well over 90% of final electricity demand in all countries (often near 100%).

The question arises, though, of the character of the unmet demand – how long duration are the unmet demand gaps? If one considers, for example, periods of supply in which only 10% of final energy are met by solar and wind, what are the longest such gaps? At 150% of annual generation and 3 hours of energy storage, but no extra interconnection, the longest duration gaps over which around 90% of power supply from solar and wind is unavailable, are in almost all cases at most on the order of a few hours to a few days. The only exception in the countries studied is South Korea, with a gap of around 600 hours, which is less than a month.

The most challenging situation we consider is the Fast Transition scenario (Figure S8, Panels E-J). From 2050 onward there is enough dedicated daily grid storage capacity in this scenario to store 40% of daily final electricity, or around 10 hours (daily grid batteries plus multi-day storage – the latter can of course be cycled daily as required). (Including the stock of EV batteries would take the total daily storage to around 60% of daily final electricity, but we don't rely on this, we simply note that there is potential here for large cost savings if this was used instead of dedicated grid storage.) In addition solar and wind generation reaches 100% of final electricity demand by 2040, 130% in 2050, then 139% after this.* The Tong et al.⁷² study indicates that in a pure solar, wind and battery system, batteries of either 3 or 12 hour duration, plus overgeneration of 100-150% will likely achieve grid reliability well above the 90% for all countries. However, the Fast Transition scenario is not a pure solar, wind and battery system – at least 20% of final energy is always supplied by dispatchable sources of hydropower, biopower, nuclear and P2X-fired generation (see panels F and M of Figure S9). According to these results then, the combination of storage, overgeneration, and (non-P2X) dispatchable power that we include in the Fast Transition scenario should alone ensure reliability well above current regulated requirements at all stages of the transition (99.97% in the US).

In addition to this already sufficient level of technology deployment, note that we also assume that enough P2X fuel is available to cover a whole month's worth of solar and wind's combined contribution to final electricity (which is easily enough to cover all the gaps identified in the 39 years of data used in the above study, if for some reason all the other flexibility options already included failed). Furthermore, in our model we assume a 12 year lifespan for daily grid batteries and a 10 year lifespan for EV batteries. However, performance requirements for bulk energy storage in grid applications are significantly lower than for transport applications, and old batteries will continue to function at lower performance for much longer than this, perhaps another 12 years. Therefore we anticipate very large quantities of second-hand batteries being available for grid applications, at lower costs than those in our model. Thus both our cost and reliability estimates are conservative.

For a different perspective we now compare this to the evidence presented in Blanco and Faaij⁵⁷. They review 35 different studies on energy storage size requirements. There is a wide range of different views. To reach a VRE penetration level of 80-94% of annual electricity demand, all but one of the 26 relevant studies estimated that storage size of around 0.01% to 1.6% of of annual electricity demand is required (these only considered the power system, though two also included heat coupling). For Renewable Energy Source penetrations higher than this (i.e. not just VRE), 79% of the 33 relevant studies estimate that storage size ranging from 0-4% of annual energy demand is sufficient (less than 15 days per year), rising to 6.4% (23 days) if 91% of studies are considered. By comparison, we are assuming one month of P2X grid backup, which corresponds to 8.3% of annual final electricity (note that P2X used in all other sectors can be stored essentially as long as required, so this is not a problem). Only two of the 35 studies estimate that significantly more than this is required (17 days more), but the strong limitations of these studies are described in the review paper. Despite the highly varied nature of the models, regions and technologies covered in Blanco and Faaij⁵⁷, the evidence strongly suggests that our Fast Transition scenario includes enough storage to surpass even the most pessimistic estimates.

It is important to understand that the reliability constraints modelled in Tong et al.⁷² already include long-duration energy storage and network expansion requirements for locations where a significant seasonal mismatch exists between VRE output and energy demand, and so by extension, we believe global seasonal storage requirements are met in our model. As a reality check on this claim, we estimate the scale of the demand variation challenge in isolation. Note that the only part of the energy system to exhibit significant seasonal demand variations is the buildings sector, due to heat and cooling requirements. Neither transport, nor industry, nor the energy sector exhibit significant seasonal demand variation (see Victoria et al.⁶³ for the

*Note that in our model "overgeneration" does not result in any curtailment or wasted energy since the large quantity of electricity that does not go straight to the grid or battery storage is used to produce P2X fuel.

European data). As shown in Table S4, the buildings sector (including residential and industry) currently accounts for 24% of final energy demand and 34% of useful energy. So we might expect buildings to account for around 30% of the future energy system. Then, the only highly populated regions currently subject to significant mismatch between seasonal supply and demand are North America and Europe (though there are of course some other sub-regions or countries that could be included). These regions currently account for around 15% of the global population, 31% of primary energy consumption, and (very roughly) half of global GDP (though their population, GDP, and energy consumption relative to the rest of the world is likely to decrease over the coming decades as developing countries are generally growing faster in all three metrics)¹⁰². Thus, only approximately 10-20% of the global energy system (i.e. only the buildings sectors of North America, Europe, and other sub-regions) will have to cope with strong seasonal demand variation. This is a only a relatively small fraction of the global energy system.

For a second reality check on seasonal demand variation, we refer to Victoria et al.⁶³, in which a sector-coupled model of the European power system was studied. Although the technologies used and model design are quite different to ours, around 550 hours (i.e. 23 days) of long-duration thermal energy storage was required in the very low emissions system studied. This is comparable to our month's worth of backup P2X storage. However, again, our model includes one month of backup P2X storage to cover final electricity in *all* sectors of the energy system, not just heat demand, and all regions. We therefore estimate that our P2X store is around 3 times larger than that identified as necessary in the above study (and is available in all regions if required, not just those with very cold winters).

Comparing different technology strategies and levels of deployment is difficult due to the array of options and interactions and different metrics and methodologies used in the various studies. We have used evidence from both the geophysical constraints perspective and the energy model literature perspective to produce a solution that appears to satisfy all energy demand requirements. We have designed the Fast Transition scenario so that it has more than enough storage capacity to meet all system demands – it has both the levels of short-term storage, over-generation and flexible generation to meet reliability standards from the geophysical perspective, and the level of long-term storage required to cover long lulls in variable renewable energy output. No further grid expansion, interconnection, or LDES additions are required.

The Fast Transition scenario has around 38 EJ of P2X fuel available for electricity generation in 2050 (producing over 5000 TWh of electricity), rising to 66 EJ in 2070 (producing over 9000 TWh of electricity). This is in addition to 100-150 TWh of short term grid storage, plus an equivalent amount of EV battery packs. For comparison Heide et al.¹⁰³ and Rasmussen et al.¹⁰⁴ estimated storage needs in Europe of around 480 TWh and 320 TWh respectively; several other studies estimated regional storage needs in the hundreds of TWh also; and Pleßmann et al.¹⁰⁵ estimated global storage needs of around 7180 TWh. Thus our method produces quantities of storage technologies that comfortably surpass estimates found in the literature, and our lack of precision in specifying the exact operational details of the electricity system in our model is more than compensated for by the massive oversupply of technologies able to meet all electricity and energy demands.[†]

3.5 Electricity adjustment for hydrogen

We assume all P2X fuels are made by electrolysis. The conversion efficiency of electrolysis is currently around 60% for PEM and 75% for SOEC, and the expected future values are around 70% and 90% respectively²⁶. We assume a constant value of 70%, so that around 1.43 EJ of electricity is required to make every exajoule (LHV) of P2X fuels in the model. The total quantity of electricity generation (from energy inputs) required by the system is the sum of final electricity from non-P2X sources, and electricity used for electrolysis:

$$Elec_t = \underbrace{EF_t^{4, total}}_{\text{Final electricity}} - \underbrace{\theta(\psi^9 + \psi^{10})EF_t^{4, total} \min\{\Pi_t^{P2X grid}, 1\}}_{\text{Final electricity from P2X fuels}} + \underbrace{\frac{1}{0.7} \times P2X_t}_{\text{Electricity for all P2X fuels}}. \quad (25)$$

3.6 Electrolyzers

Next we calculate the quantity of electrolyzers required to produce all the P2X fuels. Electrolyzers are quantified in terms of their nameplate power capacity – the rate at which they can take electrical energy and produce hydrogen. They are usually described in terms of kW (i.e. kWh/hour) but we use units of EJ/hour

[†]The Slow Transition scenarios rely far less on non-dispatchable sources of generation, while still having significant storage capacity, so easily meet very high reliability standards.

in order to align with the units used for all other technologies. We assume that the fleet of electrolyzers installed has on average a 50% utilisation rate, so they operate for half of all hours each year.[‡] As above, we assume electrolyzer efficiency of 70%. Therefore, in order to produce the quantity $P2X_t$ EJ of P2X in year t (given in Eq. 24), the total capacity of electrolyzers available on the system must be $\frac{1}{0.5} \times \frac{1}{0.7} \times P2X_t$ EJ/year. Electrolyzers are assigned technology index 13 in the model, so, in units of EJ/hour we have

$$q_t^{13} = \frac{1}{24 \times 365 \times 0.5 \times 0.7} P2X_t \quad (26)$$

This is the installed capacity of electrolyzers implied by the quantity of P2X fuels demanded in the scenario. Note that there are no scenario-dependent parameters involved in this calculation, it is determined fully by the quantity of P2X fuels and the capacity factor assumption.

3.7 Electricity networks

Electricity conversion, storage and generation technologies are represented explicitly in our model, via electrolyzers (S.I. 3.3, S.I. 6.13), two types of battery (S.I. 3.1, S.I. 3.2, S.I. 6.12), and seven different types of generators. These all depend critically on the existence of well-functioning electricity networks. We use the terms “electricity networks” and “power grid” (or just “grid”) interchangeably, and use them to mean transmission lines, distribution lines, substations, control systems, and any associated infrastructure, but *not* generators.

We model power grid costs in a slightly different way to all other technologies in the model. This is because suitably detailed historical data on electricity network costs are not available, so we are unable to make probabilistic cost forecasts as with other technologies. Instead, we use a simple strategy that allows us to make best use of the variables and data already in the model, as well as the limited power grid cost data that is available: we model the power grid as a system with two distinct capabilities that represent the main differences across scenarios: *size* and *flexibility*. (This approach is similar to that used in the IEA’s World Energy Model¹³ and the REMIND integrated assessment model¹⁰⁶.)

First, the grid must be able to deal with the sheer quantity of electricity in use in the economy – we call this the size component. This can be thought of as the power lines, transformers and various pieces of physical equipment we are used to seeing in the grid today (but as we have already said, excluding generators). This component reflects not only the distance the grid spans, but also the power capacity of the system (i.e. the “thickness of the wires”), and scales roughly with the total amount of electricity generated by the system each year.

Second, the grid must be able to balance supply and demand reliably – we call this the flexibility component. This can be thought of as the control systems and many other technologies that provide flexibility to the grid (on both the supply and demand sides), ensuring adequate resources are in the right locations at the right times. This component will vary dramatically between different scenarios, depending on the level of variable renewable energy (VRE) sources on the system.

We use two different methods to model the two components across scenarios. Since batteries, P2X production, and P2X grid backup generation are already included in the model, we use these to model the flexibility component. To model the size component, we estimate the costs required to scale up the grid to cope with more electricity in the economy, assuming all VRE integration issues are covered independently. As with all our modelling choices, this strategy is chosen for its simplicity, and because it allows us to use longer, more reliable data sets to calibrate as much of the model as possible.

3.7.1 Grid flexibility

We first describe in more detail how we model the flexibility component. The power grid allows electricity to be delivered on demand whenever and wherever it is required, constantly matching supply and demand, at all times and in all locations of the grid. Historically this balance has been maintained primarily by two methods: grid connections, which allow power to be transmitted and smoothed over space; and dispatchable electricity, which allows power to be shifted over time, by generating electricity on demand from energy stored in fuels and hydroelectric dams. However, dispatchable electricity as it existed historically was just one

[‡]We pick this central value here for parsimony. However, note that electrolyzers would likely be placed in locations with the cheapest, most reliable combinations of wind and solar, which may allow higher utilisation rates. In addition, there will likely be so much battery storage on the system that load-shifting over hours and days will be possible, again allowing for utilisation rates above 50%, so this assumption may well be pessimistic.

particular, *combined* method of performing both energy storage and electricity generation.[§] Now that both nondispatchable electricity sources (wind and solar) and standalone electricity storage solutions (batteries and electrolyzers) are available, there are many more possibilities for how to organise the grid. Since these technologies have never existed at scale though, it is not known yet what the best configuration will be.[¶] Several additional tools for managing the power grid also exist, as described in S.I. 3, but we focus on grid extensions and energy storage here, as these are likely to be the most important elements of grid management, and are sufficient to form a simple upper bound on the costs of VRE integration requirements in all scenarios.

Grid extensions, upgrades and interconnections between distant locations allow solar and wind electricity to flow from sunny and windy regions to dark and windless regions as it is generated. Energy storage allows solar and wind electricity to be stored and used later. While it is clear that storage and grid extensions are to some extent substitutes, the exact extents, given different resources and demand characteristics, are not yet known.^{||} Many studies have investigated how to incorporate nondispatchable generation sources into the grid at large scale^{77,108}. These studies have consistently concluded that storage is a more expensive solution than extending grids to cover larger regions to make use of greater resource diversity. As with all energy models though, this conclusion follows directly from assumptions about future energy technology costs, which, as shown in Figure 3, have often been highly inaccurate. Despite this current consensus, our strategy for modelling power grid flexibility costs is based around ensuring that very large quantities of energy storage are available whenever high levels of nondispatchable generation are present in a scenario.

We choose this storage-focused approach for two reasons. First, high quality historical data for power grid extension and interconnection technology costs are not available, so we are unable to make empirically-grounded cost forecasts as we are with other technologies (in contrast, we have much better data for storage technologies, so we can reliably forecast their costs). Second, since bulk storage is believed to be far more expensive than grid extensions, our resulting system cost estimates should be upper bounds on the cost of high renewables scenarios, since cheaper, grid-extension-based (or hybrid) solutions will likely be possible.

The details of how much battery and electrolyzer capacity are assumed in high VRE scenarios are shown in S.I. 3. In summary though, the Fast Transition scenario has enough utility-scale storage capacity to store around 7 hours of global final electricity each day, to deal with the diurnal cycle of solar power. It also has an additional 5 hours of global final electricity of storage per day in the form of EV batteries, which may be charged flexibly, as and when resources allow. Finally, enough excess electricity is generated to make P2X fuels for all hard-to-electrify applications, and to provide backup power for all electricity grids globally for one month (in terms of final electricity delivered), in case there are long lulls in solar and wind generation.

The costs of all these technologies are already included in the total system cost calculation. To get a sense of their scale, using a discount rate of 1.4%, and including costs through 2070, the median discounted expenditures on all grid batteries included in the Fast Transition scenario sum to 16.3 tn\$(2020) and the discounted expenditures on electrolyzers required to make P2X fuel total 3.4 tn\$(2020). These sum to a discounted value of 19.7 tn\$(2020) for flexibility-providing technologies in the Fast Transition scenario. In contrast, the No Transition scenario has none of these extra technologies (and incurs no extra expenditures), because it maintains the current power generation mix, so all generation is dispatchable. Thus, via the explicit representation of several key storage technologies that facilitate VRE integration, our model covers the flexibility component of power grid costs in all scenarios.

3.7.2 Grid size

To deal with the size component of power grid costs, we estimate how much it would cost to scale up the current, dispatchable-generation-focused power grid to serve the different levels of electricity generation present in different scenarios. (Scenarios are shown in detail in S.I. 4.) The baseline, to which we compare all

[§]For example, coal-fired electricity generation consists of a large heap of coal, which performs the energy storage function, and coal combustion, which performs the electricity generation function. The coal power system inherently includes both elements, but they can now also be done separately, e.g. via PV and batteries.

[¶]Distributed solar, storage and micro-grids are likely to be a prominent part of the most cost-effective solution. For instance, recent grid modelling efforts in the US have suggested that a configuration with significant local distributed solar and storage could save the US almost half a trillion dollars in system costs by 2050¹⁰⁷.

^{||}For example, consider a system with only PV and wind generators. In a large region with very little wind, grid connections are not a good substitute for storage, because no matter how many power lines we build, we can not access overnight wind power from far away locations to meet overnight electricity demand. The only option available for overnight electricity would be PV plus batteries. But in a large region with lots of highly dispersed wind, we can build power lines to access overnight wind power, so that both PV-plus-batteries, and PV-plus-wind-plus-grid are technically feasible. In this case grid extensions and storage are to some extent substitutes.

scenarios for this process, is the Historical Mix scenario, because the historical power grid cost data that we use applies only to this generation mix – any extra grid costs associated with increased levels of electrification of the energy system, relative to this baseline, must be accounted for in the model. Both Fast and Slow Transition scenarios have much higher levels of electrification, so require substantial extra grid investments, but the No Transition scenario is very similar to the Historical Mix, so the cost differences are minimal.

In the Historical Mix scenario, electricity supply continues to grow steadily to meet a growing global demand. New grids will be built for the many parts of the world that are currently under-served, and routine network upgrades and replacements will be performed as outdated infrastructure is retired everywhere. Investments in upgrading and extending grids are therefore required in similar proportion, relative to the economy, as they were for past investments. In the Historical Mix scenario, final electricity and total electricity generation both scale up from 26.9 PWh/yr in 2021 to 71.1 PWh/yr in 2070, a factor of 2.6 (Figure S6, panel E). In contrast, the Fast Transition has a much faster increase in solar and wind generation, which is used both to provide much more final electricity (either directly or via batteries) and to convert energy to P2X fuel for storage and other uses. In this scenario final electricity in 2070 is 133.2 PWh/yr, which is 5.0 times the 2021 value (i.e. 1.9 times the 2070 Historical Mix value) (Figure S8, panel E). Total generation (from energy inputs) in 2070, which also includes the electricity used to make P2X fuels, is 208.6 PWh/yr, which is 7.6 times the 2021 value (i.e. 2.9 times the corresponding Historical Mix value).

The Fast Transition scenario therefore represents an expansion of final electricity by a factor of 1.9, and an expansion of total generation by a factor of 2.9, *relative to the Historical Mix*. Now, in terms of circuit miles, 90% of the power grid is the distribution system, the function of which is to deliver final electricity to end-users (as opposed to the 10% that is the transmission system, see below for further details). Hence final electricity is most likely the best way to measure the scale-up in power grid infrastructure required. However, in our spirit of making conservative assumptions, we use the larger total generation scale-factor as the basis for estimating grid infrastructure expansion, and how much this will cost.** We now describe how we estimate how much the increase in grid size would cost in each scenario.

First consider the Historical Mix scenario as a baseline. The IEA’s World Energy Investment reports^{109,110} provide data on global annual investments in electricity networks from 2012-2021. The average value over this period was around 280 bn\$(2020)/yr. On average 69% of investment was in the distribution system and 31% was in the transmission system. These values remained fairly stable over the period considered (they are all within 10% of the mean). To model grid costs in the Historical Mix scenario we assume that annual investments simply scale in proportion with total electricity generation. In other words, we take the annual investment per unit of electricity generation to be a constant,

$$c_{\text{grid}} = \frac{280 \text{ bn}\$(2020)/\text{yr}}{26.9 \text{ PWh}/\text{yr}} = 10.4 \text{ bn}\$(2020)/\text{PWh}. \quad (27)$$

Then, with total generation $Elec_t^{\text{Historical Mix}}$ (as defined in Eq. 25 above), annual grid investments at time t are given by

$$Grid_inv_t^{\text{Historical Mix}} = c_{\text{grid}} \times Elec_t^{\text{Historical Mix}}. \quad (28)$$

Hence we model annual investments increasing, from 280 bn\$(2020)/yr, by a factor of 2.64 by 2070, to 739 bn\$(2020)/yr. These annual investments represent the baseline costs of replacing, upgrading and extending power grids in the Historical Mix scenario between now and 2070.

To estimate the extra expenditures that would be required in the Fast, Slow, and No Transition scenarios, over and above those in the Historical Mix scenario, we must consider how the grids would differ. In these scenarios, the same annual investments will be made on building, extending and improving grids for under-served nations as their economies develop, and replacing old grid infrastructure in developed economies. The primary difference between these grids (especially the Fast/Slow Transition grids) and the Historical Mix grid is that during the replacement and extension processes, that will happen anyway, the new infrastructure will be built with greater capacity, in order to transfer more net energy, to supply the increased electrification of

**It seems likely that using total electricity generation to estimate the required power grid infrastructure scale-up will result in an overestimation of costs. This is because P2X production on the scale seen in the Fast Transition scenario will require large industrial production facilities, which will most likely be co-located and optimised with large generating capacity, and perhaps also batteries. These facilities will not involve the complex, high-density network of distribution lines, substations, and land and building constraints that characterise the delivery of final electricity to the economy more generally. Instead they will likely require fewer, simpler, transmission lines. As explained in more detail below though, around 90% of power lines and 80% of grid expenditures are at the distribution level. Without these expenditures, costs are likely to be much lower. Hence, the quantity of final electricity delivered probably gives a better indication of grid investments required.

the economy. The number of end users will be approximately the same, they will simply receive more energy through upgraded wires.

There are two main ways to increase the capacity of the grid: install more lines (known as increasing the number of circuits), or install higher voltage lines. The latter allows higher power flow through the lines, resulting in higher net energy transfer. To estimate how much these alternatives cost, we must consider the transmission and distribution (T&D) grids separately. Distribution lines are low voltage (around 2.5kV to 34.5kV)¹¹¹, relatively cheap, and numerous. Transmission lines are medium, high, or extra high voltage (around 34.5kV to 1000kV)¹¹¹, more expensive, and far less numerous. There are about 10 million km of distribution lines in the European Union, which constitute 97% of all power lines there¹¹². In the US there are around 6.3 million miles of distribution lines, and 700,000 circuit miles of transmission lines^{113,114}. Transmission systems typically cost around 5 to 10 times the cost of distribution systems, though the specific details of each project (terrain, land costs, material requirements etc.) have a large impact on final costs^{115,116}.

We must now estimate how much it costs to increase the capacity of the transmission and distribution parts of the grid separately. The distribution component is straightforward. As stated in Short¹¹⁷, the cost of a 34.5kV distribution network is around 1.05-1.25 times the cost of a 12.5kV distribution network. Thus we can install distribution grid equipment with around three times the voltage (and hence capacity) for 1.15 times the cost on average (i.e. the midpoint of the range given).

Increasing grid capacity at the transmission level is a little more costly. To estimate this we use cost tables published by the Midcontinent Independent System Operator¹¹⁶ and the Western Electricity Coordinating Council¹¹⁸. These tables describe the approximate costs per mile of constructing new transmission lines of various types and voltages, in various locations in the US. For example, MISO¹¹⁶ states that in Illinois, to build a single circuit transmission line at a voltage of 69kV, 115kV, 138kV, 161kV, 230kV, 345kV, or 500kV, costs approximately 1.4M\$/mile, 1.5M\$/mile, 1.5M\$/mile, 1.6M\$/mile, 1.7M\$/mile, 2.8M\$/mile, or 2.9M\$/mile, respectively. Thus it is possible to build a 138kV line instead of a 69kV line (i.e. double the voltage) for just 1.07 times the cost in this case. Similarly the voltage may be tripled, to 230kV, for just 1.21 times the cost of a 69kV line. (Similar tables are given for upgrading or building new substations, but the cost ratios are very similar to those for transmission lines, so we focus only on the latter here.)

To summarise these tables, MISO¹¹⁶ estimates that installing double circuit transmission lines (i.e. doubling the number of lines along which power can flow) costs around 1.40-1.75 times the cost of single circuit transmission lines (the variation is due to different voltages, locations, terrain etc.). Doubling the transmission voltage costs around 1.15-1.80 times the base cost. Tripling the transmission voltage costs around 1.20-2.10 times the base cost. Similarly, WECC¹¹⁸ estimates that the circuit-doubling and voltage-doubling cost scale factors are around 1.6 and 2.0, respectively, though this report only considers voltages at the higher, more expensive end of the range (unlike the MISO report, which considers lower voltages too)^{††}. Averaging the MISO and WECC estimates for circuit-doubling and voltage-doubling yields a cost factor of 1.66. Therefore we can install transmission grid equipment with around twice the capacity for around 1.66 times the cost.

As described above, MISO¹¹⁶ estimates that tripling the voltage costs on average only a little more than doubling it (a factor of 1.65, compared to 1.48). However, this factor of 1.65 is very close to the *average* capacity doubling factor of 1.66 (since circuit-doubling is cheaper than voltage-doubling). But we expect tripling to cost more than doubling, so we need to increase this factor. To do this we observe that the average WECC doubling factors are slightly higher than the average MISO doubling factors (1.80 vs. 1.53). Hence we take the average tripling-to-doubling ratio from the MISO data ($1.65/1.48 = 1.12$), and apply it to the WECC-MISO-averaged doubling factor of 1.66. This results in an estimated tripling factor of 1.86. Therefore we estimate that we can install transmission grid equipment with around three times the capacity for around 1.86 times the cost, on average.

Now, since distribution systems are much cheaper but more abundant than transmission systems, we must apply the different T&D scaling factors to each component separately. As stated above, 69% of annual investment on electricity networks is spent at the distribution level, and 31% at the transmission level. Therefore, in order to triple the network capacity of any potential infrastructure investments, we apply the factor 1.86 to transmission expenditures, and the factor 1.15 to distribution expenditures, yielding a total T&D cost scaling factor of 1.37 (i.e. $0.69 \times 1.15 + 0.31 \times 1.86$). This gives

$$c_{\text{grid}}^{\text{triple capacity}} = 1.37 \times c_{\text{grid}} = 14.9 \text{ bn}\$(2020)/\text{PWh}. \quad (29)$$

Finally then, to upgrade the grid from the Historical Mix scenario grid to those required in the Fast/Slow/No Transition scenarios, we assume that as routine extension and replacements are made, some fraction of these

^{††}The WECC report does not consider voltages low enough to allow tripling, so we are unable to provide this figure.

are built at triple capacity (either through circuit or voltage upgrades), at 1.37 times the cost. This fraction is determined by comparing the total electricity generation in the different scenarios. Since total generation in the Fast Transition scenario is nearly three times that in the Historical Mix scenario by 2070, almost all of the investments will need to be made at the higher level. The fraction of the Historical Mix grid that must be tripled is given by

$$\frac{1}{2} \left(Elec_t^{\text{scenario}} - Elec_t^{\text{Historical Mix}} \right). \quad (30)$$

This is shown graphically in Figure S4.




Figure S4: Scaling up the power grid to cope with the quantity of total generation in the Fast Transition, as compared to the Historical Mix, does not simply require more power lines, but also higher capacity power lines. This is because the number of end-users in both cases is the same, they simply receive more electricity through upgraded wires. The shaded black region shows the fraction of Historical Mix power grid investments that must be made at the higher, capacity tripling level ($c_{\text{grid}}^{\text{triple capacity}}$) in the Fast Transition scenario, in order to increase the capacity sufficiently. The dotted lines show both the “+” and the “-” parts of the mathematical formula in order to emphasise geometrically the construction of the fraction of grid investments that must be tripled – the three segments are of equal size, and the lowest (shaded) segment is fully contained within the Historical Mix total generation.

To obtain the annual investment in electricity networks required in any scenario, we simply sum the single-capacity investment component and the triple-capacity component,

$$Grid_inv_t^{\text{scenario}} = c_{\text{grid}} \times \left(Elec_t^{\text{Historical Mix}} - \frac{1}{2} \left(Elec_t^{\text{scenario}} - Elec_t^{\text{Historical Mix}} \right) \right) \quad (31)$$

$$+ c_{\text{grid}}^{\text{triple capacity}} \times \left(\frac{1}{2} \left(Elec_t^{\text{scenario}} - Elec_t^{\text{Historical Mix}} \right) \right) \quad (32)$$

These annual investments are included in the total system cost calculation, and in this way the model covers the size component of power grid costs. It should be noted that the No Transition scenario is similar enough to the the Historical Mix scenario, in terms of total electricity generation, that the extra grid investments required for No Transition are small. Both Fast and Slow Transition scenarios require much higher additional grid costs. These results are shown in Figure S44. Including grid investment costs in the total system cost calculation in this manner makes it straightforward to calculate total Fast and Slow Transition scenario costs relative to No Transition costs, as is the ultimate aim.

There are four reasons why this cost estimate is a conservative upper bound on the cost of upgrading the grid to deliver much more electricity in the Fast and Slow Transition scenarios. First, upgrading existing infrastructure is usually cheaper than installing new equipment, because many of the costs have been paid already, e.g. land, acquisition, permitting etc.¹¹⁹. Therefore, since we have used new installed system costs here, many of the upgrades required to existing systems will be cheaper than these estimates. Second, energy storage investments allow grid investments to be deferred, because they help even out power flow, preventing large peaks from overloading equipment. A very large number of such storage investments are already

accounted for in the model in the Fast and Slow Transition scenarios, so fewer grid investments may be required than are estimated here. Third, many advanced economies are experiencing a legacy of over-capacity following recent increases in energy efficiency. In some locations, existing under-utilized grid capacity may be sufficient to deliver almost all of the required increase in total electricity generation (providing new storage technologies are added to smooth out power flow and avoid the large demand-driven peaks and troughs observed in the current system). Finally, as mentioned already at the start of this section, we are using total electricity generation (from energy inputs) to scale grid investments, whereas final electricity, which is much lower, may be a more appropriate scaling factor, due to the way the grid would be laid out in high VRE scenarios.

In summary, the model includes both the costs of providing power grid *flexibility*, via storage technology deployments, and power grid *size*, via annual grid investments that scale up in accordance with total electricity generation. As a reality check, observe Figure S46, which shows median annual investment results produced by our model. The fraction of investments devoted to electricity networks in our transition scenarios is of the order of 10%. This is approximately the same as the fraction observed in the high renewables, highly system-coupled European energy system scenarios modelled using the high-resolution power systems analysis model PyPSA-Eur-Sec-30 in Brown et al.³⁷ (see their Figure 6 in particular). This gives us confidence that our strategy for modelling grid investments is adequate.

3.8 Technology lifespan and annual additions

In order to calculate the costs of many of the technologies in the model we need to know not only the total annual quantity of production of the technology, q_t^i , but also the vintage of the underlying capital stock. This is because for these technologies (all electricity generation and storage technologies) “quantities” are defined in terms of annual production, while costs are accounted for in terms of annuitized (levelized) costs, which persist for the entire lifespan of the underlying capital stock. In these cases we must keep track of the annual additions, and use costs associated with the *year of installation* when calculating total costs of production.

For example, technology $i = 10$ is “electricity generated by solar PV panels”, and q_t^{10} is the quantity of this technology produced in year t . But, by the definition of LCOE, the same solar panel produces energy at the same LCOE for its whole lifespan. So to calculate the total cost of producing q_t^{10} we must add up the individual costs associated with panels of every vintage in the installed capacity base (each with a different LCOE), as well any new annual additions.

Let the lifespan of technology i be L^i periods. We assume that the underlying capital stock functions perfectly for L^i periods then is removed from the system. Let $Q_{t,\tau}^i$ be the quantity of technology i produced in year t by capital stock that was installed in year $\tau \leq t$. For simplicity we assume that the first year’s production, q_1^i , relies on capital stock that was installed linearly over the previous L^i periods. Hence if there were no further additions, production would just decay linearly over the next L^i periods, representing constant annual removals of the underlying capital stock: we would have $q_t^i = \max\{0, q_1^i(1 - (t-1)/L^i)\}$. However, accounting for new additions on top of this initial decaying production gives

$$q_t^i = \sum_{\tau=1}^t Q_{t,\tau}^i \quad (33)$$

$$Q_{t,\tau}^i = \begin{cases} \max\{0, q_1^i(1 - \frac{t-1}{L^i})\} & \text{for } \tau = 1 \\ \max\{0, q_t^i - \sum_{s=1}^{t-1} Q_{t,s}^i\} & \text{for } \tau = t \text{ and } t > 1 \\ Q_{t-1,\tau}^i & \text{for } 1 < \tau < t \text{ and } 1 < t \leq L^i \\ \max\{0, Q_{t-1,\tau}^i - Q_{t-L^i,\tau}^i\} & \text{for } 1 < \tau < t \text{ and } t > L^i \end{cases} \quad (34)$$

Eq. 33 says that production in year t is the sum of production relying on capital stock installed in all past years, plus the current year. Working down the rows of Eq. 34, the top row says that any production relying on capital stock from year 1 decays linearly for L^i years and is zero thereafter. The second row says that in year t , the amount of production relying on brand new capital stock, $Q_{t,t}^i$, is the difference between the total required production, q_t^i , and all production relying on previous years’ capital stock – these are the annual additions. The third row says that at all times t before one full lifespan of the technology has passed ($t \leq L^i$) the underlying capital stock just rolls over, so the resulting production this year is the same as last year. The final row says that for any times *after* the first lifespan has passed, underlying capital stock of age L^i is removed. Figure S5 shows how the process works.

We have thus determined q_t^i for all i and t .




Figure S5: Annual solar PV generation broken down by PV system vintage in a slow deployment scenario, left (initial annual growth 10%), and a fast deployment scenario, right (initial annual growth 30%). Colours show the median forecast LCOE in the deployment year. For any given year, the block of energy at the *top* of the stack corresponds to energy generated by PV systems installed that year. Blocks of energy lower down the stack are generated by older PV systems with higher LCOEs, and the block at the base of the stack is in its last year of life before being removed. Fast deployment allows access to large quantities of very cheap energy in the 2030s, but slow deployment makes this unlikely before the 2050s, as experience is not gained fast enough.

4 Scenarios

There are two critical questions when considering future scenarios: how fast will different technologies grow, and for how long? Based on the history of deployment of a few energy technologies, Kramer and Haigh¹²⁰ argued that there are “societal laws” that necessarily slow exponential growth once an energy technology passes 1% of primary energy supply. However, very little historical data was presented in that paper, with no supporting analysis or validation. The main argument for why growth must slow was that only 2-4% of existing capital stock needs replacing each year, and “industry will only consider early retirement of the existing capital stock if the total cost of the new technology (capital and operating costs) falls below the operating cost of the old”¹²⁰. First, this does not explain the 1% threshold, but more importantly, it says nothing about how growth proceeds when total annuitized costs *do* drop below existing technology operating costs. Wind, solar and batteries differ from previous energy technologies because their costs have so far dropped exponentially at significant rates. They are also highly modular and can be built very quickly. This makes them fundamentally different from previous energy technologies such as nuclear power. While it is clear that storage is needed for wind and solar to go far beyond 50% of the electricity supply, there is still substantial room for growth before that point is reached (which at current rates should occur in about 12 years). Regarding the replacement rate, as we argue here, when coupled with growth and economically motivated asset stranding, replacement rates are sufficient for the growth we envisage.

We construct five scenarios, listed below. In all cases useful energy in each end-use sector grows at 2% until 2100. (We vary this assumption in Section 7.3.) All technology growth rates are less than or equal to their recent observed rates, apart from two minor exceptions.* Hence we do not need to assume technologies grow faster than they are currently growing, we simply vary the length of time for which this growth continues. In fact, growth rates for solar, wind, batteries and electrolyzers over the last twenty years have often been higher than the rates we use, so it is possible to construct faster transition scenarios while still being broadly consistent with recent historical trends, but in keeping with our conservative approach we avoid this. Sections 4.1-4.5 show the scenarios in more detail.

- **No Transition:** This scenario represents a situation where fossil fuel use continues to grow in line with historical trends for another decade then begins to plateau. As the fossil fuel plateau proceeds, solar, wind and the electrification of transport finally take off at a meaningful scale to impact the full system technology mix. For this scenario to occur, solar and wind only grow at around 10% per year, much slower than their historical averages, and together reach approximately 50% of electricity generation in around 2060. We use this as a baseline to compare to other transition scenarios, as it is similar in character to many of the “Current Policies” or “Stated Policies” scenarios available in the literature since 2018. Indeed, the initial growth rates of all major technologies in this scenario are taken directly from such scenarios found in various IEA WEO reports, from 2018-2021. The main difference between this scenario and comparable ones in the literature is that since we assume a constant 2% growth in total useful energy, something has to make up the growth after fossil fuels have plateaued, and for this we use solar and wind (because they are cheapest). While this means that our No Transition scenario is not directly comparable with other “Stated Policy” type scenarios (which have plateauing useful energy consumption), this approach is required for the apples-to-apples cost comparison we perform in this work. (Note that prior to 2018, typical “Stated Policies” scenarios were referred to as “Business-As-Usual” and were closer to our “Historical Mix” scenario shown below.) See Figures S6 and S7 and Table S9.
- **Fast Transition:** In this scenario, we suppose that clean energy and storage technologies grow at rates consistent with historical data until they become dominant, and then slow down to the system-wide growth rate of 2% by 2040. This sees solar and wind grow to account for substantial shares of the energy mix by 2040. Fossil fuels are quickly displaced from all sectors, mostly through electrification, but also by replacement with electrolytic P2X fuels, especially in industry and transport. By 2040, the installed capacity of electrolyzers, short-term batteries and multi-day storage have grown to a scale able to contribute significantly to energy storage requirements. Oil consumption falls dramatically due to the rapid electrification of transport – this is one of the first big transitions to occur, because while EVs grow at a similar rate to the other storage technologies, they start from a much higher level so reach

*These are: nuclear energy in the Slow Nuclear Transition, since this case requires a sharp increase in nuclear energy; and gas electricity in the Fast Transition scenario, but this transient increase is just an artifact of the scenario construction process, and could in any case be dealt with easily via increased utilisation rates of the existing CCGT fleet.

mass scale sooner. After five decades of sustained 2% growth of transport services, final energy in the sector is still below current levels due to the much higher efficiency of EVs.

By 2040, enough short-term battery capacity exists to store and shift 20% of all solar and wind electricity generated each day, and enough multi-day storage exists to store 10% of daily solar and wind generation. Enough P2X fuel is produced to cover one month's worth of solar and wind's contributions to the power grid in case there is an extended lull in output due to the natural variation of these sources. By 2040 90% of electricity is generated from zero carbon sources and 81% of final energy is provided by zero carbon sources, rising to 91% in 2050 and 97% in 2060. See Figures S8 and S9 and Table S10.

- **Slow Transition:** In this scenario, the transition to a clean energy system is generally slow until 2050 and then gathers pace. We assume solar and wind grow at around half their recent historical rates and maintain these levels for the next few decades. This has the effect of gradually greening the electricity supply by 2050 (to around 80%, with the rest coming from gas). Significant fuel switching in transport does not occur and final energy remains well above current levels. Although decarbonisation in transport and industry begins immediately in this scenario, progress is so slow that only by mid-century does this begin to show significant results. With the slower growth of solar and wind, far less electricity needs to be stored and there are still significant amount of dispatchable fossil fuel generators available. This also means that much less P2X fuel is required for grid backup – only one week's cover of solar and wind's contribution is provided. Less P2X fuel is needed in transport and industry because fossil fuels continue to be used. By 2070 roughly 88% of final energy comes from clean energy sources. See Figures S10 and S11 and Table S11.
- **Slow Nuclear Transition:** This scenario is almost the same as the Slow Transition, except that nuclear grows at 8% per year and eventually takes a very large share of electricity generation. This results in an electricity system with a far lower proportion of solar and wind, so there are no grid batteries and no P2X fuels grid backup. By 2070 roughly 88% of final energy comes from clean energy sources. See Figures S12 and S13 and Table S12.
- **Historical Mix:** In this scenario the current mix of technologies stays fixed, while the whole system expands at 2% per year. See Figures S14 and S15 and Table S13.

4.1 No Transition




Figure S6: No Transition scenario details (continued on next page). Historical growth rates shown in panel D are 5-year moving averages.

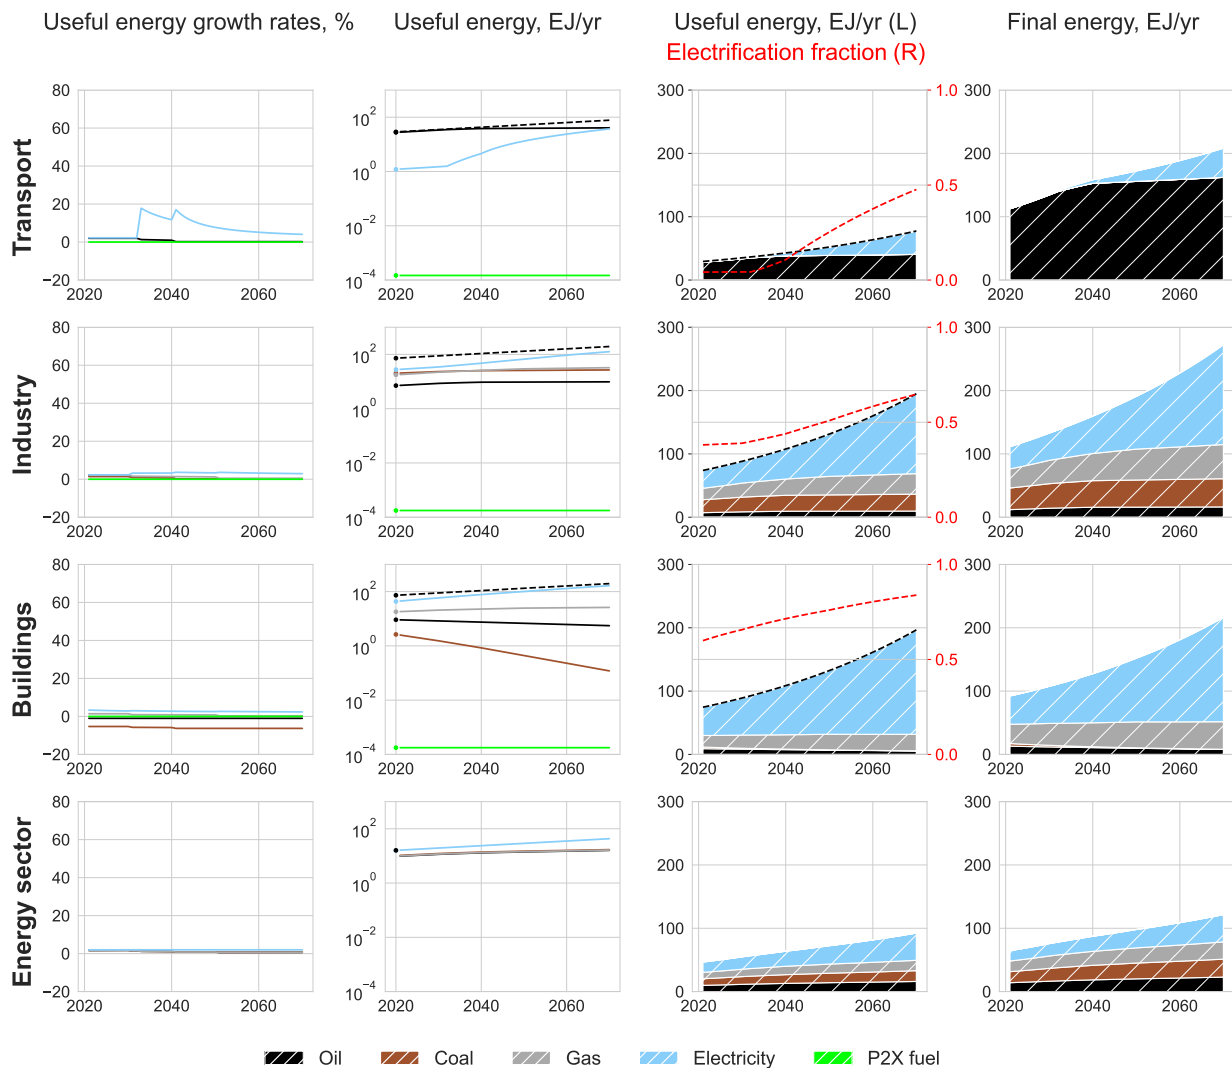



Figure S7: Top: Continuation of Figure S6. Main: No Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

		Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3		
Sector	Energy carrier						Max. % of useful energy	
Transport	Oil	1.99	0.2	12	20	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	0	0	-	-	-	-	
Industry	Oil	1.9	0.1	10	20	-	-	
	Coal	1.4	0.2	10	20	-	-	
	Gas	2.3	0.5	10	30	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	0	0	-	-	-	-	
Buildings	Oil	-1.0	-1.0	10	30	-	-	
	Coal	-5.3	-6.3	10	20	-	-	
	Gas	1.4	0.3	10	30	-	-	
	Electricity (slack)	0	0	-	-	-	-	
	P2X fuels	0	0	-	-	-	-	
Electricity	Technology						Max. % of generation	
Generation	Coal	1.5	1.0	10	30	-	-	
	Gas	2.5	1.5	10	30	-	-	
	Nuclear	1.5	1.4	10	20	-	-	
	Hydropower	1.1	1.1	10	20	-	-	
	Biopower (slack)	-	-	-	-	-	-	
	Wind	8	5	10	30	1	28	
	Solar PV	12	5	10	30	1	28	
Storage	EV batteries	10						
	Daily grid batteries	0	~0% of daily solar+wind generation stored ($\beta_{ST} = 0.0001$)					
	Multi-day storage	0	~0% of 3-day solar+wind generation stored ($\beta_{LT} = 0.0001$)					
	P2X fuels		No solar+wind grid backup ($\theta = 0$)					

Table S9: No Transition construction parameters. Initial and long term growth rates are taken from IEA WEO reports from the years 2018, 2019, 2020 and 2021.

4.2 Fast Transition

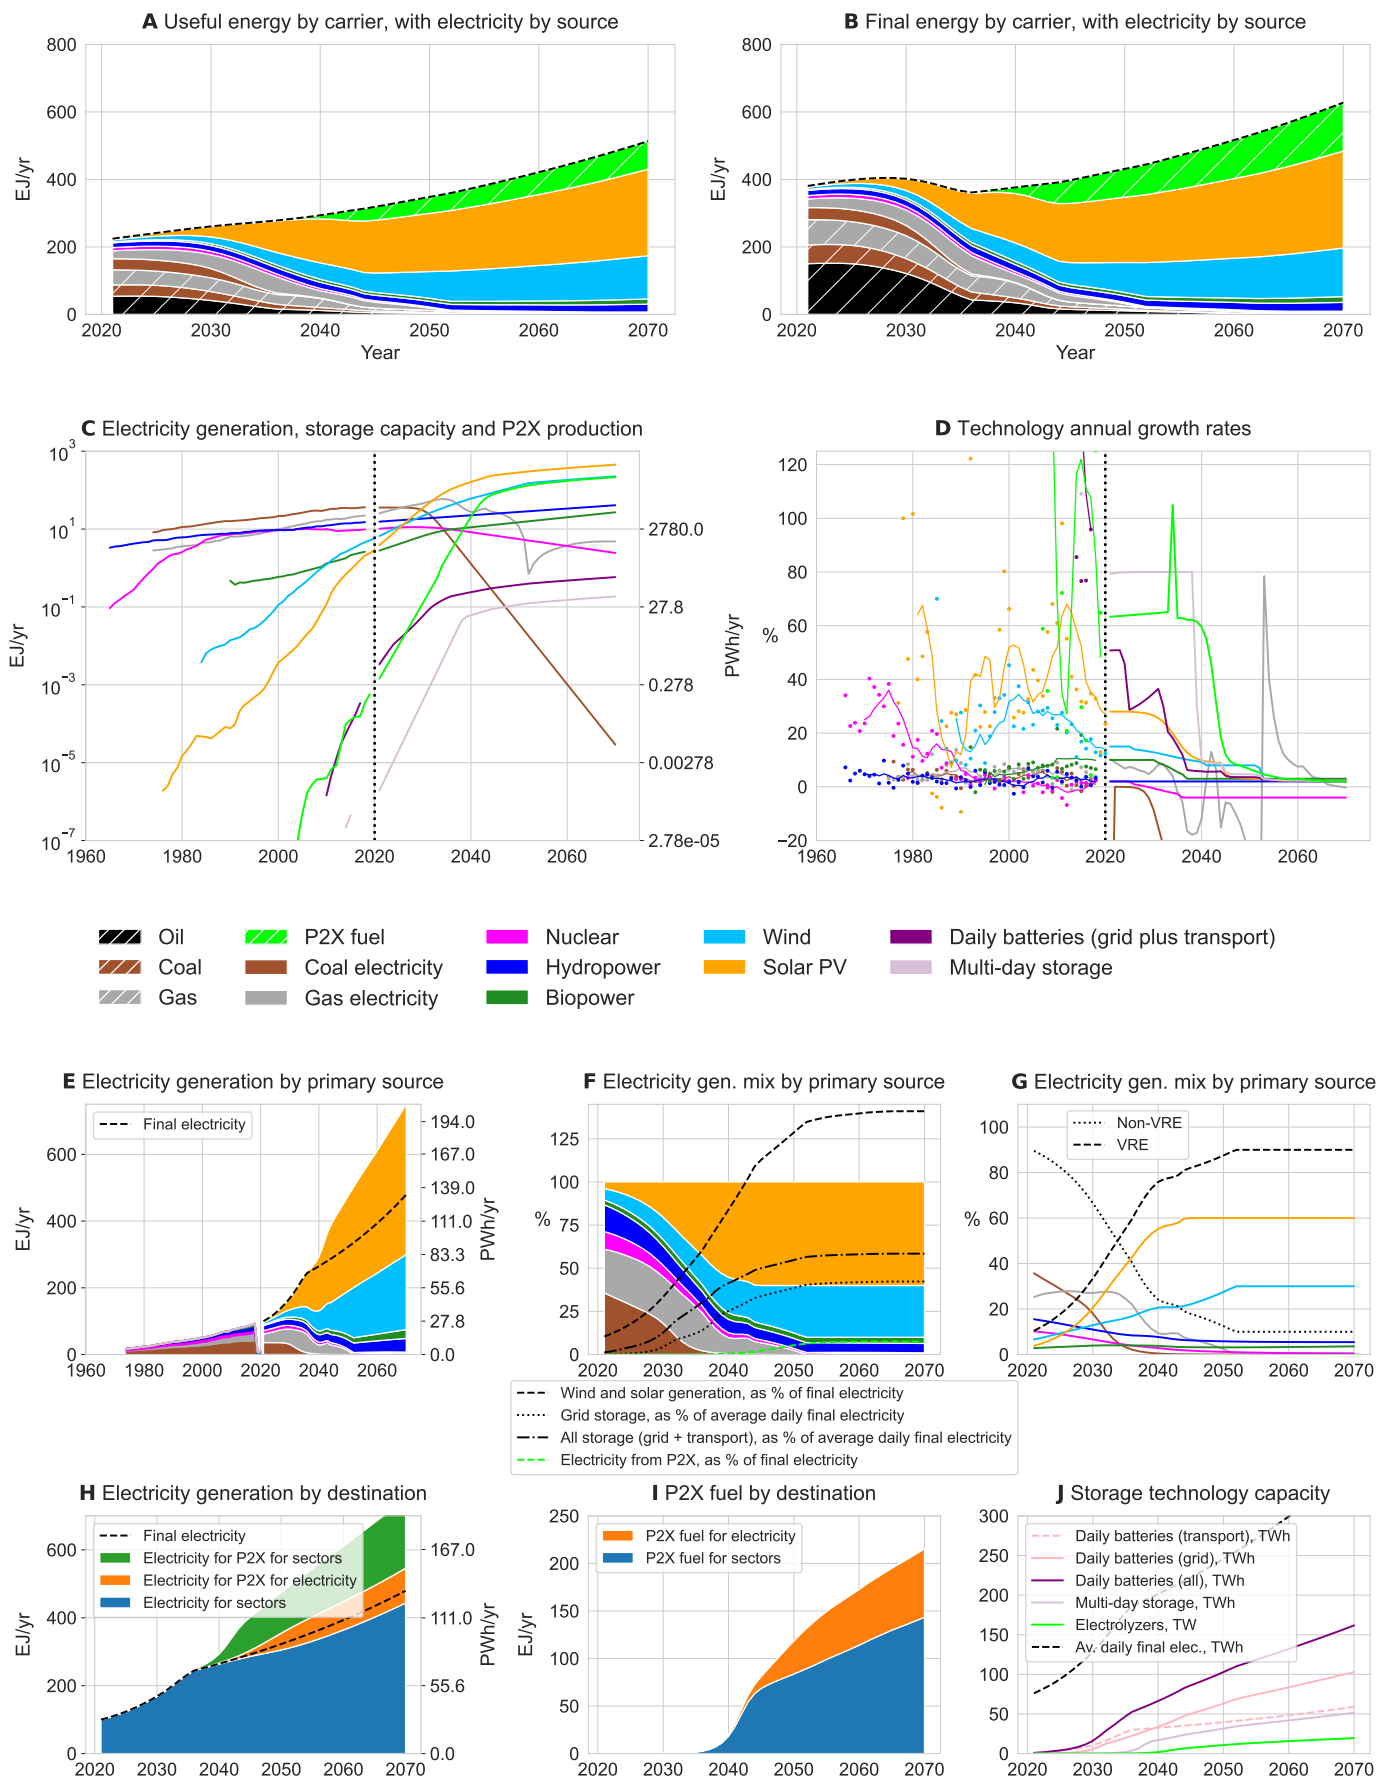


Figure S8: Fast Transition scenario details (continued on next page).





Figure S9: Top: Continuation of Figure S8. Main: Fast Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

		Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3		
Sector	Energy carrier						Max. % of useful energy	
Transport	Oil (slack)	-	-	-	-	-	-	
	Electricity	25	5	8	20	1	80	
	P2X fuels	60	4	13	31	1	20	
Industry	Oil	0	-15	2	10	-	-	
	Coal	0	-12	2	14	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Electricity	7	4	8	32	1	75	
	P2X fuels	70	3	2	41	1	24.5	
Buildings	Oil	0	-16	1	8	-	-	
	Coal	0	-16	1	8	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Electricity	4	8	4	30	1	89.5	
	P2X fuels	60	4	13	31	1	10	
Electricity	Technology						Max. % of generation	
Generation	Coal	0	-30	2	18	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Nuclear	2	-4	5	15	-	-	
	Hydropower	2	2	10	30	60	5.5	
	Biopower	10	3	10	16	1	4.2	
	Wind	15	8	5	20	1	30	
	Solar PV	28	9	5	22	1	60	
Storage	EV batteries	50						
	Daily grid batteries	60	20% of daily solar+wind generation stored ($\beta_{ST} = 0.2$)					
	Multi-day storage	80	10% of 3-day solar+wind generation stored ($\beta_{LT} = 0.1$)					
	P2X fuels		1 month of solar+wind grid backup ($\theta = 1/12$, $t_{P2X\ grid} = 40$)					

Table S10: Fast Transition construction parameters.

4.3 Slow Transition

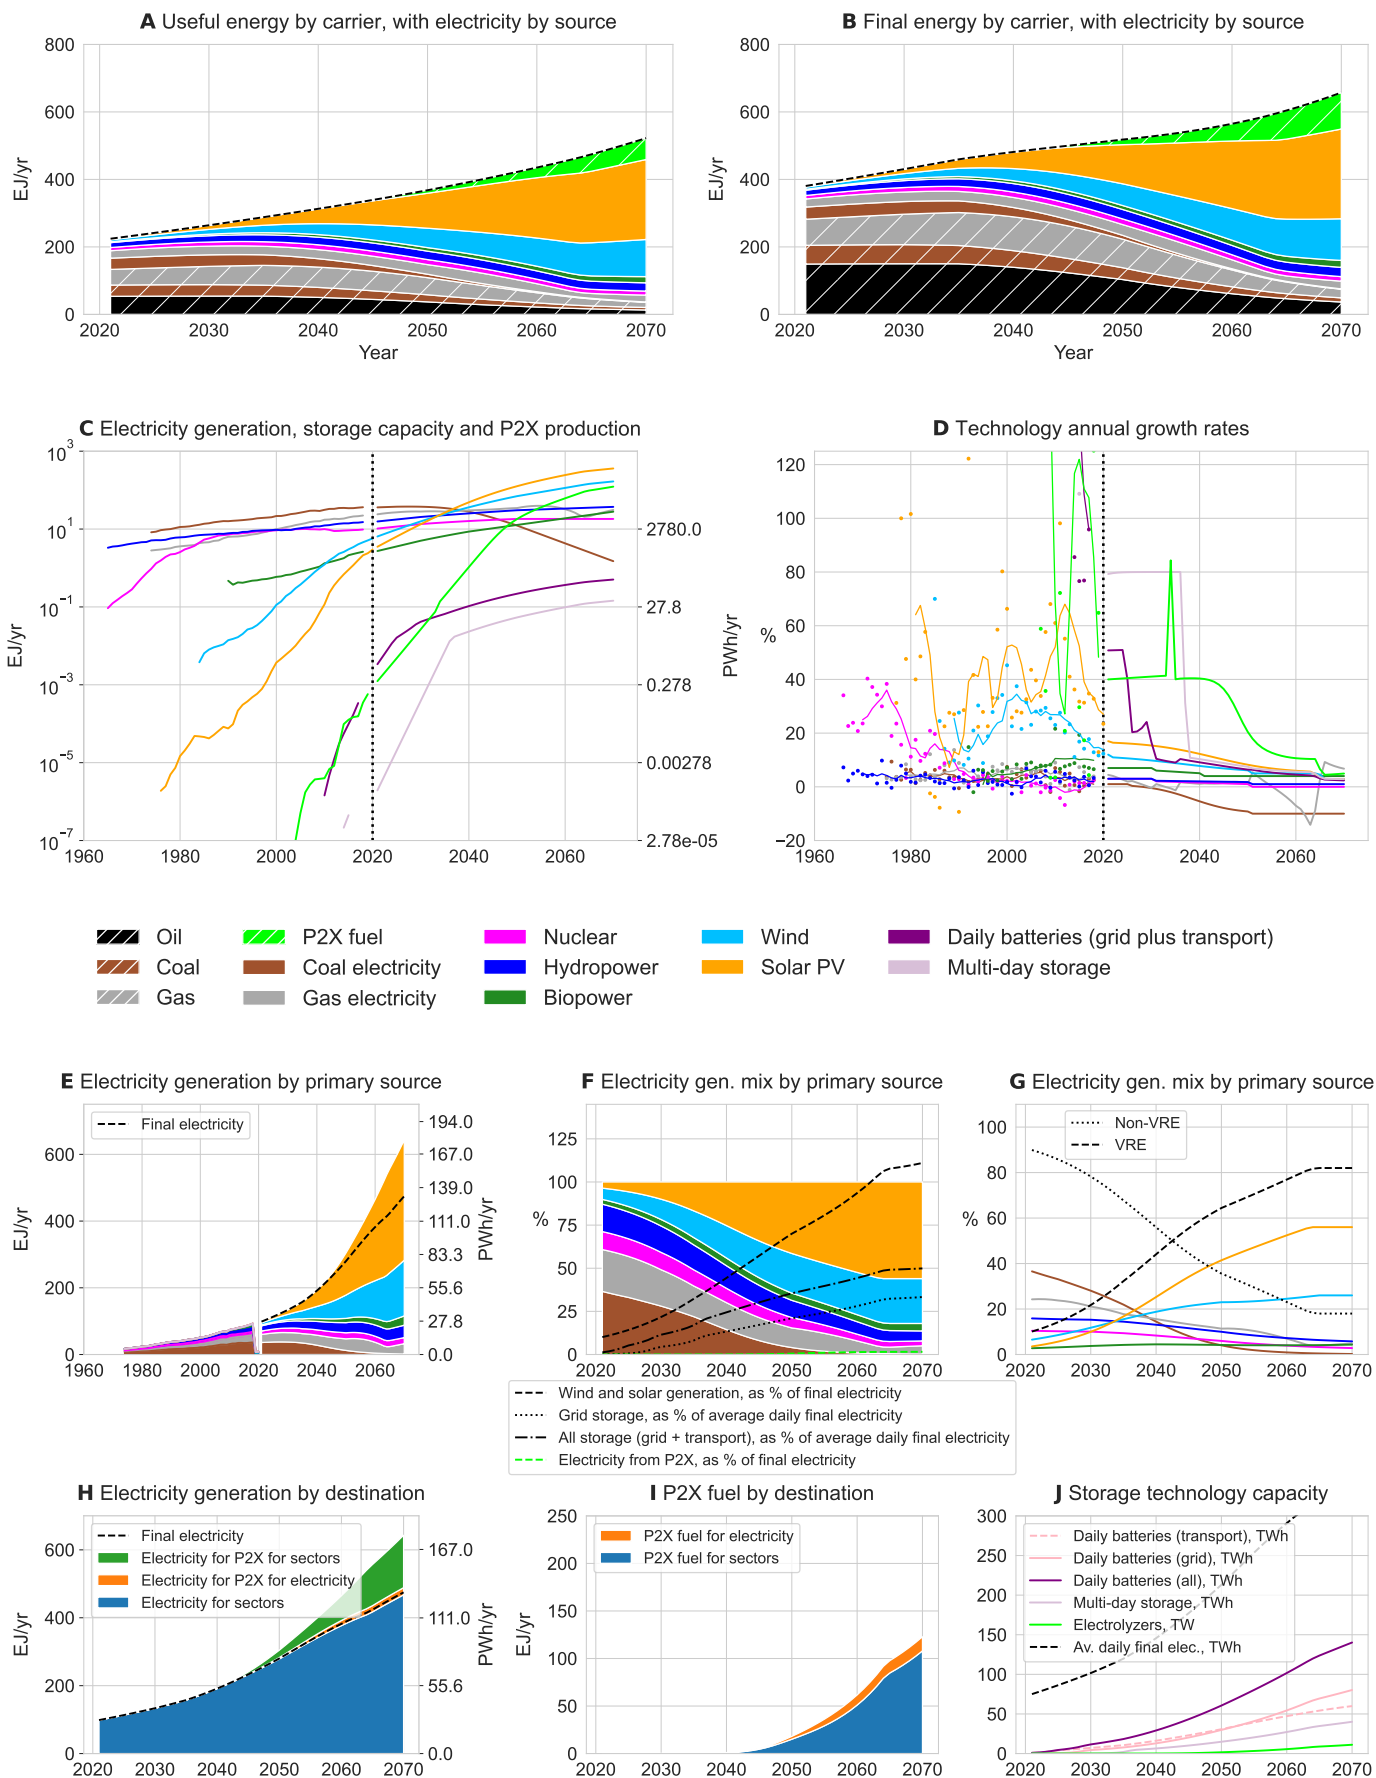


Figure S10: Slow Transition scenario details (continued on next page).





Figure S11: Top: Continuation of Figure S10. Main: Slow Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

		Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3		
Sector	Energy carrier						Max. % of useful energy	
Transport	Oil	0	-5	15	30	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	35	12	12	42	1	18	
Industry	Oil	0	-6	15	30	-	-	
	Coal	0	-6	10	30	-	-	
	Gas	2	-6	15	40	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	45	12	12	42	1	20	
Buildings	Oil	0	-6	15	30	-	-	
	Coal	0	-5	15	30	-	-	
	Gas	2	-9	15	40	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	40	10	12	42	1	10	
Electricity	Technology						Max. % of generation	
Generation	Coal	1	-10	5	30	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Nuclear	3	0	10	30	-	-	
	Hydropower	3	1	10	30	-	-	
	Biopower	7	4	10	20	1	12	
	Wind	12	5	1	30	1	26	
	Solar PV	17	5	1	43	1	56	
Storage	EV batteries	50						
	Daily grid batteries	60	20% of daily solar+wind generation stored ($\beta_{ST} = 0.2$)					
	Multi-day storage	80	10% of 3-day solar+wind generation stored ($\beta_{LT} = 0.1$)					
	P2X fuels		1 week of solar+wind grid backup ($\theta = 1/52, t_{P2X\ grid} = 50$)					

Table S11: Slow Transition construction parameters.

4.4 Slow Nuclear Transition




Figure S12: Slow Nuclear Transition scenario details (continued on next page).





Figure S13: Top: Continuation of Figure S12. Main: Slow Nuclear Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

		Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3		
Sector	Energy carrier						Max. % of useful energy	
Transport	Oil	0	-5	15	30	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	35	12	12	42	1	18	
Industry	Oil	0	-6	15	30	-	-	
	Coal	0	-6	10	30	-	-	
	Gas	2	-6	15	40	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	45	12	12	42	1	20	
Buildings	Oil	0	-6	15	30	-	-	
	Coal	0	-5	15	30	-	-	
	Gas	2	-9	15	40	-	-	
	Electricity (slack)	-	-	-	-	-	-	
	P2X fuels	40	10	12	42	1	10	
Electricity	Technology						Max. % of generation	
Generation	Coal	1	-10	5	30	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Nuclear	8	4	20	45	1	40	
	Hydropower	3	1	10	30	-	-	
	Biopower	3	3	10	20	-	-	
	Wind	11	5	1	30	1	22	
	Solar PV	15	5	1	40	1	28	
Storage	EV batteries	50						
	Daily grid batteries	60	~0% of daily solar+wind generation stored ($\beta_{ST} = 0.0001$)					
	Multi-day storage	80	~0% of 3-day solar+wind generation stored ($\beta_{LT} = 0.0001$)					
	P2X fuels		No solar+wind grid backup ($\theta = 0$)					

Table S12: Slow Nuclear Transition construction parameters.

4.5 Historical Mix

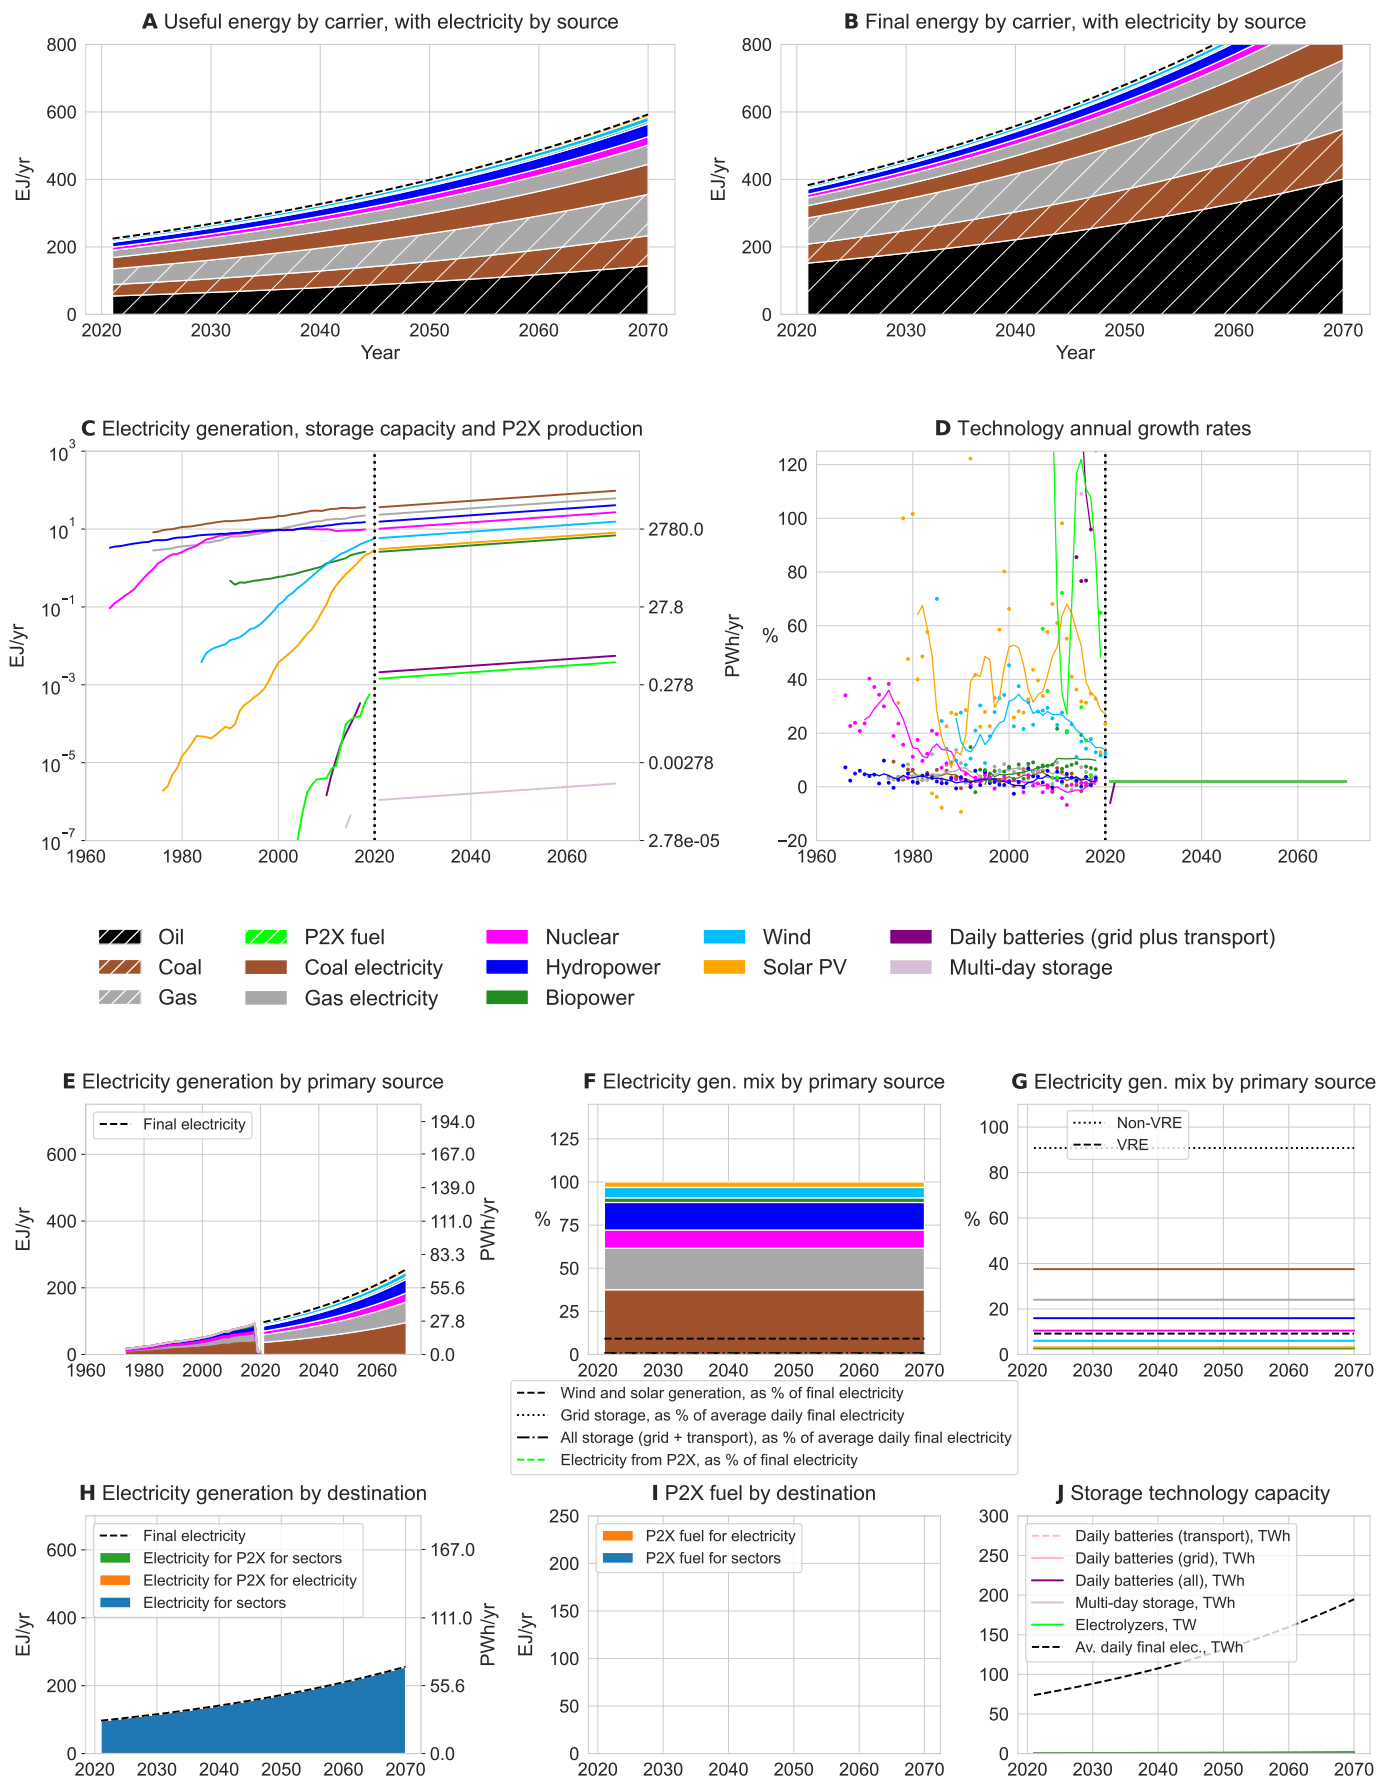


Figure S14: Historical Mix scenario details (continued on next page).

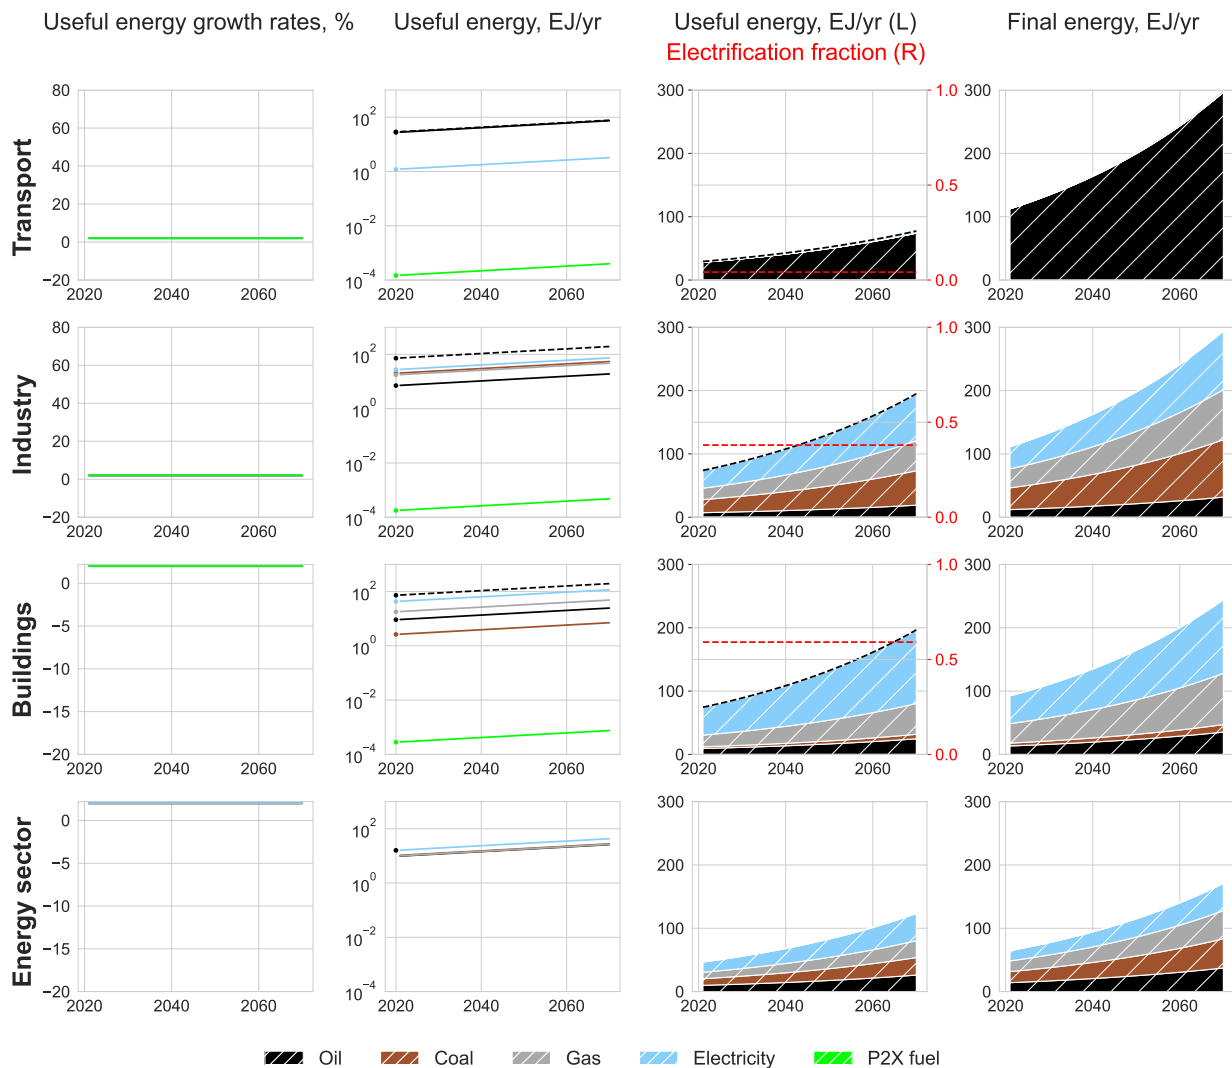



Figure S15: Top: Continuation of Figure S14. Main: Historical Mix carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

		Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3		
Sector	Energy carrier						Max. % of useful energy	
Transport	Oil (slack)	-	-	-	-	-	-	
	Electricity	2	2	-	-	-	-	
	P2X fuels	2	2	-	-	-	-	
Industry	Oil	2	2	-	-	-	-	
	Coal	2	2	-	-	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Electricity	2	2	-	-	-	-	
	P2X fuels	2	2	-	-	-	-	
Buildings	Oil	2	2	-	-	-	-	
	Coal	2	2	-	-	-	-	
	Gas (slack)	-	-	-	-	-	-	
	Electricity	2	2	-	-	-	-	
	P2X fuels	2	2	-	-	-	-	
Electricity	Technology						Max. % of generation	
Generation	Coal	2	2	-	-	-	-	
	Gas	2	2	-	-	-	-	
	Nuclear	2	2	-	-	-	-	
	Hydropower	2	2	-	-	-	-	
	Biopower (slack)	-	-	-	-	-	-	
	Wind	2	2	-	-	-	-	
	Solar PV	2	2	-	-	-	-	
Storage	EV batteries	2						
	Daily grid batteries	2	~0% of daily solar+wind generation stored ($\beta_{ST} = 0.0001$)					
	Multi-day storage	2	~0% of 3-day solar+wind generation stored ($\beta_{LT} = 0.0001$)					
	P2X fuels		No solar+wind grid backup ($\theta = 0$)					

Table S13: Historical Mix construction parameters.

4.6 Comparison with AR6 scenarios

To provide some context for the five scenarios constructed, we compare some key features with those of the thousands of scenarios included in the Sixth Assessment Report (AR6) of the Intergovernmental Panel on Climate Change¹²¹. The comparisons are shown in the ten panels of Figure S16. We show all AR6 scenarios contained in the AR6 Scenarios Database (World, v1.0) (i.e. not just the “vetted” scenarios), which was downloaded from the AR6 scenario explorer¹²².

Panels A and B of Figure S16 show that the deployment of solar energy in the Fast Transition scenario is at the top of the range of trajectories seen in AR6 scenarios (though not outside it). Panels C and D show that wind deployment in all our scenarios is closer to the middle of the AR6 range (notwithstanding the many AR6 trajectories well below current deployment levels).

Panels E-I show different aspects of final energy in the various scenarios. Comparing scenarios is always difficult because they consist of so many different variables, which vary in different ways over time, but the five panels shown here provide a good overview of some of the most important differences. They should be interpreted as a group (along with the solar and wind panels already shown), so that the relationships between different variables make most sense. Note that in this section (and here alone), we use different terminology than the rest of the paper. In the IPCC database, to the best of our knowledge, “final energy” refers to *end-use energy only*, that is, energy carriers used in end-use sectors, but not the energy sector. We therefore use this terminology here, and in each case, the “final energy” presented for our scenarios is only the final energy used in end-use sectors – we have removed the energy sector self-consumption component.

Panel E shows final electricity. Since all our scenarios (except Historical Mix) have high levels of electrification *and* constantly increasing useful energy, final electricity is high compared to most AR6 scenarios (which tend to project useful energy growth stalling by mid-century). Fast Transition is right at the top of the range, although the two Slow Transition scenarios are also high. Note that at around 2035 in the Fast Transition, although final electricity switches to a lower growth rate, electricity generation as a whole continues to grow strongly (see solar and wind in panels B and D), and excess generation is used for P2X production. (Note that in Section 7.3 we vary the exogenous useful energy growth rate – effectively scaling down final electricity so that our scenarios lie nearer the middle of the AR6 range – and show that our results are robust to this choice.)

Panel F shows final liquids (i.e. oil, biofuels, and P2X fuels such as green ammonia). This is almost a mirror image of panel E, reflecting how transport electrification occurs in parallel with sharply falling oil demand, while identical, increasing levels of transport services are maintained. Since our model calculates total P2X production, but does not specify which fuels are produced, to estimate the amount of final liquids and gases in our scenarios we have split the total final P2X fuel quantity equally between final liquids (e.g. ammonia, methanol) and final gases (e.g. hydrogen, methane). Thus the quantities plotted here for our scenarios are: [final oil + 0.5 × final P2X] in panel F, and [final gas + 0.5 × final P2X] in panel H. These trajectories are indicative, and the P2X shares can be modified up or down depending on which specific fuels are required. (Note that these plots show only *end-use* P2X, and that in the Fast Transition scenario further P2X is produced for power grid backup in the energy sector. This feature is not easy to capture with any of the metrics available in the AR6 data set.) Panel F reflects how in the Fast Transition scenario oil demand falls rapidly as transport electrification occurs, but then around 2035, as oil continues to fall towards zero, P2X production picks up and provides substitute fuels for use in non-electrifiable applications.*

Panel G shows the coal component of final energy solids. Our scenario trajectories behave as expected. The other significant final solids component not shown here is biomass, but this is not included in our model, so we do not include it in the comparison.

Panel H shows final gases (i.e. methane, syngas, and P2X fuels such as green hydrogen). Note that in the Fast Transition scenario fossil gas decreases gradually until 2035, similar to final liquids above, following which it continues to fall towards zero as it is replaced by P2X fuel. Thus, even in the Fast Transition scenario, which reaches near zero emissions, final gas ends up near the middle of the range of the AR6 scenarios, due to high P2X production.

Panel I shows total system final energy. Our scenarios have an intentionally low starting point because we

*The kink around 2035 reflects this fuel switch, and is an artifact of our scenario construction method that would not be expected to occur in reality. Panel H shows that the artifact is not present for final gas. This is because final gas is spread more evenly over end-use sectors, so there is a slight averaging effect in our model as P2X gas substitutes for fossil gas at different rates in different sectors. In contrast, oil is concentrated entirely in the transport sector, so there is no such averaging effect. This could easily be addressed with a more complex scenario construction method. Alternatively, if instead of the 50-50 split of P2X fuels between liquids and gases that we have assumed here, we constructed a P2X liquid-gas split that varied over time, it would be quite simple to remove this modelling artifact. For simplicity though we avoid this, and emphasize that these trajectories are merely indicative.

omitted several components of the energy system from our model. The largest omitted component by far is traditional biomass, at around 26 EJ (see Table S2), which makes up most of the difference seen on the plot. Despite the low starting point though, both Fast and Slow Transition scenarios end up with high final energy – above the mid-point of the AR6 scenario ensemble. Hence the final energy *growth* of our scenarios is likely at the very top of the range of growth in all AR6 scenarios, if not above it. (Similar to with final electricity above, in Section 7.3 we run a test of scaling scenario final energy up and down, and demonstrate that our main conclusions are unchanged.) The large dip in final energy up to 2035 in the Fast Transition scenario is caused by large efficiency gains due to the electrification of transport.

Panel J shows energy system CO₂ emissions. Again the starting point for our scenarios is a little low, due to the omitted model components, but the scenario trajectories behave as expected.

Overall the ten panels show that our scenarios, which are based on an energy services growth rate similar to rates observed historically, provide high levels of energy services relative to the AR6 ensemble. Given this, our scenarios appear to cover a fairly representative range of the space of scenarios, covering wide deployment ranges of all major technologies. Our approach of exogenously holding useful energy growth constant across scenarios shows clearly how structural changes (i.e. energy carrier substitutions), and their rates, can yield energy systems with identical levels of energy services available for powering the economy, which nevertheless have entirely different environmental characteristics (greenhouse gas emissions, air quality etc.). These comparison plots help show how such structural differences impact the position of our scenarios within the ensemble of AR6 scenarios, in terms of a number of key variables.

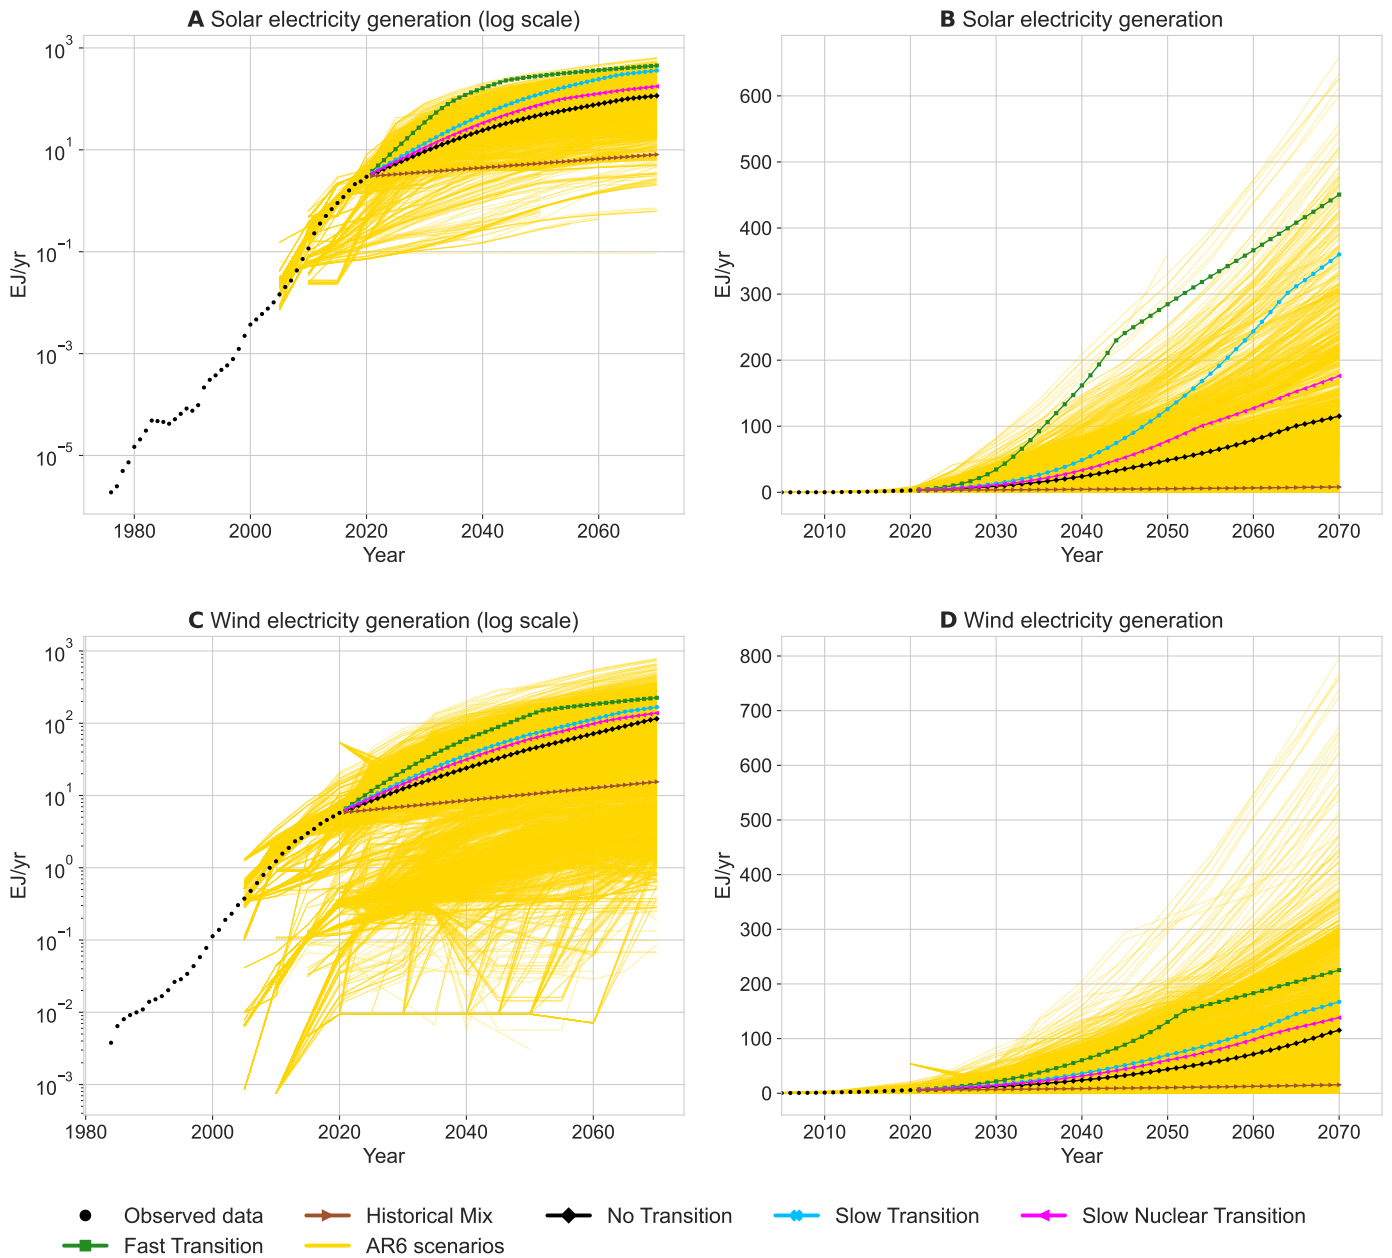


Figure S16: Scenario comparisons on log and non-log scales, respectively, for solar (A and B), and wind (C and D). (Figure continues on next page.)




Figure S16: Scenario comparisons for several key scenario metrics. See text for full description.

5 Technology cost models

In this section we discuss our calculation of the present discounted cost of each scenario, relative to No Transition. Our approach involves three steps: (1) determine the quantities of each energy technology produced in each scenario; (2) make cost forecasts for each technology based on historical data; (3) calculate present discounted costs associated the quantities of component technologies for each scenario and sum these up to arrive at a total present discounted cost estimate.

When making technology forecasts in step 2 we use two different cost models depending on whether or not a technology's costs have historically remained approximately flat over time. Fossil fuel based technologies (including electricity from fossil fuels) have tended to remain fairly stationary and are in the first category while all other technologies are in the second category. As we discuss below, we use an autoregressive (AR(1)) time series model for the first category and a Wright's law model for the second category. (In S.I. 7.5 we show the effects of using Moore's law instead of Wright's law.) (In addition to these two probabilistic cost models, due to a lack of suitable data we model annual investment costs for technology $i = 14$ (the electricity network) using a much simpler, deterministic method, as described in S.I. 3.7.)

An extensive literature exists on the experience curve concept and its place within the technological change and innovation literature. For an overview of its origins, theory and limitations see e.g. Thompson¹²³, Arrow¹²⁴, Mazzola and McCardle¹²⁵, Nordhaus¹²⁶ and Way et al.¹²⁷. For details of applications in the field of energy technologies see e.g. Neij¹²⁸, Grübler et al.⁴⁷, Rubin et al.¹²⁹, Samadi¹³⁰, Wene¹³¹, Ferioli et al.¹³² and Wilson et al.¹³³. For a discussion of caveats and subtleties around experience curves see S.I. 8.3.

5.1 AR(1) process

While highly unpredictable in the short-term, real fossil fuel prices have remained remarkably constant – within less than one order of magnitude – for a century or more. This is in contrast to solar PV, for example, which has dropped in price by a factor of more than 3,000 in sixty years. There have been dramatic technological improvements in fossil fuel industries at all stages of the production process, but resource extraction has become more difficult. For reasons that are not well-understood, these two effects seem to have roughly cancelled each other for more than a century, throughout the lifetime of the industry.

Pindyck¹³⁴ analysed time series properties of long-run fossil fuel prices and found that both long-term mean reversion and short-term stochastic fluctuations were essential characteristics to incorporate in a forecasting model. He tested a multivariate autoregressive model that allowed for fluctuations in both the level and slope of the trend line. Shafiee and Topal¹³⁵ summarises the many attempts that have been made to improve forecasting methods for fossil fuel prices. We choose a simple AR(1) process because empirically it is as good or better than any of the known models, and we use it to model future costs of direct-use oil, coal and gas, plus coal- and gas-fired electricity.

The AR(1) model has no long-term trend. While one might argue that oil and gas show a weak long-term increasing price trend, there is no evidence for this for coal, which costs almost the same now as it did in 1880 (see Figure S19). An equally plausible model for oil and gas is that they have no long-term trend, but there were breaks in the price levels in the 1970s and 1980s, respectively, as suggested by the data. We make this assumption here, and calibrate the models for oil and gas to data from 1973 and 1985, respectively, up to the present time.

We define the AR(1) process used in this paper as follows: let c_t be the cost of a given technology at time t , let $X_t = \log(c_t)$, and assume this develops according to

$$X_t = \phi X_{t-1} + \epsilon_t + \kappa, \quad \text{with } \phi \in (0, 1], \text{ I.I.D. noise } \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \text{ and constant } \kappa \neq 0. \quad (35)$$

ϕ is the autoregression parameter and ϵ_t are periodic random shocks. This may also be written in Ornstein-Uhlenbeck form, which shows the mean-reversion of the process more clearly,

$$X_t = X_{t-1} + (1 - \phi)(\mu - X_{t-1}) + \epsilon_t, \quad \text{where } \mu = \frac{\kappa}{1 - \phi}. \quad (36)$$

$\mu = \mathbb{E}[X_t] = \frac{\kappa}{1 - \phi}$ is the equilibrium or long-term mean. σ_ϵ is the degree of volatility around the long-term mean caused by shocks and the autoregression parameter ϕ is the rate at which the shocks dissipate.

Long term behaviour is revealed by writing X_t in terms of the previous N steps, i.e. by substituting $X_{t-1}, X_{t-2}, \dots, X_{t-N}$ in Eq. 35. Then

$$X_t = \kappa \sum_{j=0}^{N-1} \phi^j + \phi^N X_{t-N} + \sum_{j=0}^{N-1} \phi^j \epsilon_{t-j}. \quad (37)$$

Setting $N = t$ gives the expression in terms of the initial condition, X_0 :

$$X_t = \kappa \sum_{j=0}^{t-1} \phi^j + \phi^t X_0 + \sum_{j=0}^{t-1} \phi^j \epsilon_{t-j}. \quad (38)$$

Thus

$$\mathbb{E}[X_t|X_0] = \kappa \sum_{j=0}^{t-1} \phi^j + \phi^t X_0 \quad (39)$$

Then using the sum of the first n terms of a geometric series, $\sum_{j=0}^{n-1} r^j = \left(\frac{1-r^n}{1-r}\right)$, we also have:

$$\mathbb{E}[X_t|X_0] = \kappa \sum_{j=0}^{t-1} \phi^j + \phi^t X_0 = \kappa \left(\frac{1-\phi^t}{1-\phi}\right) + \phi^t X_0 = \mu(1-\phi^t) + \phi^t X_0, \quad (40)$$

and hence $\mathbb{E}[X_t|X_0] \rightarrow \mu$ as $t \rightarrow \infty$. Further manipulation gives $\text{Var}(X_t|X_0) = \sigma_\epsilon^2 \frac{1-\phi^{2t}}{1-\phi^2}$. Then since $c_t = e^{X_t}$, c_t is log-normal, we have $\text{Median}(c_t) = e^{\mathbb{E}[X_t|X_0]} = e^{\mu(1-\phi^t) + \phi^t X_0}$, and hence the long-term median value is e^μ .

Using the data shown in S.I. 6 we fit the model for each fossil-fuel-based technology and obtain estimates of ϕ , σ_ϵ and κ , which are denoted $\hat{\phi}$, $\hat{\sigma}_\epsilon$ and $\hat{\kappa}$.

5.2 Wright's law

To model future costs of non-fossil-fuel-based technologies we use the first-difference stochastic Wright's law (FDWL) analysed by Lafond et al.². This model uses historical experience curve data to produce distributional experience curve cost forecasts, i.e. forecast distributions that depend explicitly on future scenarios. The model assumes that there are two sources of uncertainty contributing towards the errors we make in our forecasts: first, periodic random shocks due to the inherent unpredictability of the future, and second, our inability to infer the true experience exponent of any given technology from the limited historical data available. Thus there are two variables to estimate in the model for each technology, the experience exponent ω , and the variance, σ_η^2 , of the distribution of the periodic noise shocks. Calibrating the model on historical data provides us with estimates of these distributions.

Let c_t be the cost and z_t the cumulative production of a given technology at time t . The basic model is:

$$\log c_t - \log c_{t-1} = -\omega (\log z_t - \log z_{t-1}) + \eta_t \quad \text{with IID } \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \quad (41)$$

where ω is the technology's experience exponent and η_t are periodic IID noise shocks, with technology-specific variance σ_η^2 . In order to reflect evidence that there were positive autocorrelations in the data¹³⁶, Lafond et al.² used a version of the model in which the noise shocks are autocorrelated,

$$\log c_t - \log c_{t-1} = -\omega (\log z_t - \log z_{t-1}) + u_t + \rho u_{t-1} \quad \text{with IID } u_t \sim \mathcal{N}(0, \sigma_u^2), \quad (42)$$

so we use this version also. The noise shocks are related by $\eta_t = u_t + \rho u_{t-1}$, and we therefore have

$$\sigma_u^2 = \sigma_\eta^2 / (1 + \rho^2). \quad (43)$$

Lafond et al.² calculated the autocorrelation parameter ρ for 50+ technologies, found the average to be $\rho^* = 0.19$, and used this value throughout their work. We use the same value here.

The parameter estimation procedure assumes $\omega \sim \mathcal{N}(\hat{\omega}, \hat{\sigma}_\omega^2)$, so when we calibrate the basic model (Eq. 41) on historical data we obtain estimates for the three values $\hat{\omega}$, $\hat{\sigma}_\omega$ and $\hat{\sigma}_\eta$. Eq. 43 is then used to compute $\hat{\sigma}_u$. To generate a single cost sample path of length T (conditional on known values of z_t), we first pick a value of ω from the distribution $\mathcal{N}(\hat{\omega}, \hat{\sigma}_\omega^2)$, then pick T successive noise shocks from the distribution $\mathcal{N}(0, \hat{\sigma}_u^2)$ and use Eq. 42 to compute c_t at each step. Repeating this sampling process a large number of times recovers the theoretical forecast distribution calculated in Lafond et al.². This process accounts for both our uncertainty in the experience exponent (due to limited historical observations), and intrinsic future uncertainty (represented by the exogenous periodic shocks).

Note that since c_t is log-normal, if the log of experience grows linearly at rate $\mu_z = \log z_t - \log z_{t-1}$, the expected cost grows as $\mathbb{E}[c_t] \sim e^{(\hat{\omega}\mu_z + \frac{\hat{\sigma}_u^2}{2})t}$. This means it always diverges as t gets large, even if $\mu_z = 0$, i.e.

even if the technology ceases to be produced. In contrast the median grows as $median[c_t] \sim e^{\hat{\omega}\mu_z t}$ and remains constant if $\mu_z = 0$. The mode grows as $mode[c_t] \sim e^{(\hat{\omega}\mu_z - \hat{\sigma}_u^2)t}$, and so goes to zero as $t \rightarrow \infty$ if $\hat{\sigma}_u^2 > \hat{\omega}\mu_z$.

Lafond et al.² showed that the forecasting accuracy (for the mean) scales as $\tau + \tau^2/m$, where τ is the forecast horizon (the time in the future for which the forecast is made) and m is the number of historical data points used to fit the parameters for the forecast. The first term corresponds to diffusion and the second term to errors caused by parameter misestimation. This means that initially the forecast errors are dominated by diffusion and the errors grow slowly, but once $\tau^2/m > \tau$ the errors are dominated by parameter estimation and grow more quickly. The crossover occurs when $\tau = m$, i.e. for a forecast horizon equal to the number of historical data points used for estimation. Thus in cases where we have 30 data points, for example, the forecast error grows reasonably slowly until about 2050 and then grows faster after that.

We are occasionally faced with data points in non-sequential years, either because the original data contains gaps or because we choose to chain together non-sequential datasets for greater historical coverage. The model validation in Lafond et al.² only considered sequential data, so technically, the forecasting performance results obtained there do not apply to non-sequential data. However, if applied sensibly to non-sequential data the differences are very small, and the forecasts are still good. We tested this by estimating the model parameters first on a full set of sequential data, then repeatedly removing data points and re-estimating the parameters; the parameters are stable throughout this process. Since the model uses first differences, calibration on non-sequential data simply means pooling the first differences from the sequential data on either side of the gap. Providing these sequences are from reliable sources, are of similar quality, and are not from overlapping time periods, there is no reason to think that increasing the amount of data we have to describe the Wright's law relationship in this way shouldn't result in better model calibration. In practise, most of the series we use have either no or very few data gaps, so the point is mostly theoretical.

In S.I. 6 we use data to obtain parameter estimates $\hat{\omega}$, $\hat{\sigma}_\omega$, and $\hat{\sigma}_\eta$ for each non-fossil-fuel-based technology.

5.3 Total costs of AR(1) and Wright's law technologies

Let the unit cost of technology i at time t be c_t^i , and the total cost of production of q_t^i be C_t^i . The two fundamentally different cost models need to be accounted for differently when calculating expenditures in the model; we consider them in turn now.

5.3.1 AR(1) model costs

For each fossil fuel based technology ($i = 1, 2, 3, 4, 5$) the cost in each period is forecast by an AR(1) model calibrated using its historical time series. The cost of technology i production at time t is

$$C_t^i = c_t^i q_t^i. \quad (44)$$

5.3.2 Wright's law model costs

For each non-fossil fuel based technology ($i = 6, \dots, 13$) we use levelized (i.e. annuitized) costs, so we need to account for the vintage of the underlying capital stock.

First consider non-fossil electricity generation ($i = 6, \dots, 10$). We take the cumulative experience variable in the FDWL model to be cumulative electricity generated. However, the new LCOE each year, c_t^i , only applies to production *additions* – all previous years' production continues to be generated at the old LCOE values. Recall from S.I. 3.8 that the quantity of technology i produced in year t from capital stock installed in year τ is $Q_{t,\tau}^i$, and the total production is $q_t^i = \sum_{\tau=1}^t Q_{t,\tau}^i$. The cost of producing all technology i in year t is therefore

$$C_t^i = \sum_{\tau=1}^t c_\tau^i Q_{t,\tau}^i. \quad (45)$$

This applies for all non-fossil electricity generation technologies.

The three remaining technologies, short-term batteries, multi-day storage and electrolyzers are slightly different because they are described in terms of installed capacity and capacity costs. The total installed quantity in year t is q_t^i , and the $Q_{t,\tau}^i$ give the breakdown by installation year. The relevant cumulative experience variable for the FDWL model is the total cumulative production of the technology, which is the initial cumulative production plus all additional capacity installations. This will be much larger than the installed capacity base at any given time due to capacity retirements. The relevant cost variable for the FDWL model is the installed system cost (i.e. capital investment cost). We assume that this cost is annuitized over the lifetime

of the technology (L^i), so that the cost in each year is only a fraction of what the installed system cost was in the year of installation. The installed system cost is annuitized over L^i periods using the capital recovery factor (CRF)

$$CRF^i = \frac{r(1+r)^{L^i}}{(1+r)^{L^i} - 1}, \quad (46)$$

where r is the interest rate, which we take to be $r = 8\%$.

Computing the total costs for medium term storage ($i = 12$) and electrolyzers ($i = 13$) is then straightforward. We assume that these technologies have only ever been, and will only ever be, used for the purpose considered in this model. Hence it is clear what cumulative production is initially, and how it increases over time as capacity is installed and eventually retired. Capacity is given in EJ and EJ/hour respectively, and costs are given in \$/GJ and \$/(GJ/hour). In year t , $Q_{t,\tau}^i$ of capacity exists that was installed in year τ , so if the installed system cost in year τ was c_τ^i , then the cost of production of $Q_{t,\tau}^i$ is $(CRF^i c_\tau^i Q_{t,\tau}^i)$. The total cost of producing all technology i , q_t^i , is

$$C_t^i = \sum_{\tau=1}^t CRF^i c_\tau^i Q_{t,\tau}^i. \quad (47)$$

Finally, consider daily-cycling batteries ($i = 11$). These are used in both the power grid and transport sectors (recall Eq. 19: $q_t^{11} = q_t^{11,grid} + q_t^{11,transport}$). The relevant cumulative experience variable in the FDWL model is the total cumulative production of batteries in both applications, q_t^{11} . However, note that historically most Li-ion batteries have been used in the electronics sector, so we must include this quantity in estimating the initial cumulative production value for the FDWL model (as well as the smaller quantity in transport and the tiny contribution from the power sector). In computing future cumulative production though, we assume that in any scenario with significant quantities of batteries, transport and grid applications will dominate the contribution from electronics by far, so we can neglect the latter. And in any scenario without significant battery quantities, the lack of batteries means that any potential inaccuracy in the cost forecast due to omitting electronics production would have negligible impact on the results anyway.

We use FDWL to calculate the installed system cost in year t and annuitize it as before. However, in computing the total scenario cost we only consider the contribution from power grid batteries, as these are the “extra” cost associated with any particular transition scenario (recall EV batteries are considered just part of routine capital stock turnover). Thus, assuming $q_t^{11,grid}$ and $q_t^{11,transport}$ are broken down by installation year as usual, the cost of production of $q_t^{11} = \sum_{\tau=1}^t Q_{t,\tau}^{11,grid} + Q_{t,\tau}^{11,transport}$ is

$$C_t^{11} = \sum_{\tau=1}^t CRF^{11,grid} c_\tau^{11} Q_{t,\tau}^{11,grid} \quad (48)$$

5.4 Total system costs

In addition to the 13 AR(1) and Wright’s law technologies, which have probabilistic cost forecasts, there are power grid capacity investments. These are simple deterministic annual investments calculated based on the quantity of electricity generation in a given scenario (see S.I. 3.7). These are added to the total costs of all other technologies. Therefore, the total cost of all production at time t in a given scenario is

$$W_t^{scenario} = \sum_{i=1}^{13} C_t^i + Grid_inv_t^{scenario} \quad (49)$$

and the net present cost of the scenario is

$$V^{scenario}(r) = \sum_{t=1}^T e^{-rt} W_t^{scenario}, \quad (50)$$

with discount rate r .

5.5 Net present cost of transition

We want to compare the likely costs associated with different scenarios, but there are various ways to do this. We take No Transition to be a baseline scenario against which to compare others. For a given scenario and discount rate r we define the net present cost of transition, $NPC^{scenario}(r)$, to be the relative cost of the scenario:

$$NPC^{scenario}(r) = V^{scenario}(r) - V^{No\ Transition}(r) \quad (51)$$

Costs are stochastic, so by generating a length- T cost sample path for each of the N technologies (by sampling the experience exponents, AR(1) noise shocks and FDWL noise shocks), and repeating the process a large number of times, we can simulate all the relevant distributions. To obtain the distribution of NPC-of-transition values for a given scenario, we calculate the net present cost of the scenario conditional upon one complete ensemble of realised parameters and shocks, then subtract the net present cost of No Transition given exactly the same ensemble, and repeat many times. Finally this allows us to calculate

$$\mathbb{E}[NPC^{scenario}(r)] \quad (52)$$

$$\text{and } median[NPC^{scenario}(r)], \quad (53)$$

which are the metrics we use to evaluate scenarios, understand likely net costs, and how they vary with discount rate and other parameter assumptions.

6 Data, calibration and technology forecasts

We use data from many sources, mostly free and openly available on the internet, but occasionally via standard university-wide subscription licenses held by the University of Oxford.

Production data comes mostly from the International Energy Agency¹² and BP's Statistical Review of World Energy¹³⁷. Cost data is much harder to find and comes from a wide variety of sources including, among others, Lazard's Levelized Cost of Energy Analyses¹³⁸, the International Renewable Energy Agency¹³⁹, the U.S. Energy Information Administration's Annual Energy Outlooks⁸⁸, Bloomberg New Energy Finance (BNEF) and Bloomberg L.P. (via Bloomberg Terminal).

For electricity generation costs, we use LCOEs (see S.I. 1.7). Due to the scarcity of LCOE data, especially for technologies other than wind and solar, we have also included some information from the EIA reporting historical LCOE *projections*. These are near-term projected LCOE values calculated, in the past, by EIA analysts, based on the evidence they had at the time. They were published in the Annual Energy Outlook (AEO) reports, though not consistently, so the picture is far from complete. The values reported were calculated based on detailed data and assumptions used in the EIA's National Energy Modeling System (NEMS) (such as capital costs, capacity factors etc.). Often, both near-term and long-term LCOE projections were reported (around 5 and 20 years respectively, though this varies). While we ignore the long-term projections, we do assume that the near-term projections were grounded firmly enough in the data of the day and have consistent enough methodology to have significant meaning (though we are aware that methodologies used in such calculations typically change over time). Strictly these values do not qualify as *observed* data, but they do give a good indication of what the actual values probably were around around the time of publication.

Finally, different organisations often report quite different LCOE values for a given technology in a given year. This is because many assumptions go in to constructing the LCOE figure, and changing these assumptions changes the result. We do our best to use the most up to date, consistent data. In some cases the variations are so wide though that we perform the whole NPC analysis using alternative parameter choices. We thus have our "main case", which uses very conservative choices of base year LCOE values and model parameters for all clean energy technologies, plus, in addition a small number of "side cases" which show how results change if we vary initial cost assumptions (or other parameters). We now provide more detail on each technology in turn.

Although cost forecasts for all fossil-fuel based technologies are made using the AR(1) model, we nevertheless plot the data for coal- and gas-fired electricity in the form of an experience curve, in addition to the time series. This is to demonstrate the qualitative difference between the experience curves for these technologies and those for non-fossil-based technologies.

For fossil fuels we make use of price data to estimate long-range trends. This is because we have price data series spanning more than a century. In contrast, the cost of fossil fuel production is highly variable by source and data is difficult to obtain. The LCOE for renewables already contains markups along the supply chain, and only omits the final markup, which is small. It is therefore a reasonable point of comparison to fossil fuel prices.

6.1 Data sources for Figure 1

Data sources for Figure 1 are shown in Table S14. This data is a representative selection of the larger data set we have collected, which is presented in detail throughout S.I. 6.2-6.13. Although in some cases cost or price data over extended time periods was only available for the US, all other comparable data we have collected, for different regions or shorter periods, indicates that this data is in line with global trends. Some US costs are higher than the global average (e.g. nuclear LCOEs), and others are lower (e.g. gas prices), but on balance we believe the overall trends depicted represent the global situation well.

To approximate useful energy we used constant conversion efficiency factors. In reality these values have changed a little over time, and will continue to do so, but our method serves as a good first order approximation. The values used are shown in Table S15, and are based on the discussion in S.I 1.8.

Technology	Metric	Years	Source
Crude oil	Price (global av.)	1861-2020	bp Statistical Review of World Energy 2021 ¹³⁷
Coal (primary)	Price (USA)	1882-1948	McNerney et al. ¹⁴⁰
		1949-2020	EIA
Gas (primary)	Price (USA)	1922-2012	EIA wellhead price
		2013-2020	EIA Henry Hub
Coal electricity	LCOE (USA)	1902-1987	McNerney et al. ¹⁴⁰
		1996-2009	EIA Annual Energy Outlook reports
Gas electricity	LCOE (USA)	2010-2020	Lazard
		1982-1989	EIA
		1990-1997	Colpier and Cornland ¹⁴¹
		1998-2008	EIA
Nuclear electricity	LCOE (USA)	2009-2020	Lazard
		1971-1996	Koomey and Hultman ¹⁴²
Wind electricity	LCOE (global av.)	2008-2020	Lazard
		1983-2020	IRENA 2021 - Renewable Power Generation Costs in 2020
Solar PV electricity	LCOE (global av.)	1957-1975	Nemet ¹⁴³ , data extracted from Figure 8
		1976-2009	Nemet ¹⁴³ , data from PCDB
Batteries (pack)	Cost (global av.)	2010-2020	IRENA 2021
		1995-2016	IEA Global EV Outlook 2018
PEM electrolyzer	Installed system cost (global av.)	2017-2020	BNEF Executive Factbook 2021
		2004-2019	Glenk et al. ¹⁴⁴
P2X fuel	Cost (global av.)	2020	IEA WEO 2020
		2004-2020	Calculated assuming a 50-50 split between solar and wind for electricity input, and utilisation rate 50%.
Crude oil	Production	1870-1964	Smil ¹⁴⁵
Coal (primary)	Production	1965-2020	BP ¹³⁷
		1800-1964	Smil ¹⁴⁵
Gas (primary)	Production	1965-2020	BP ¹³⁷
		1890-1964	Smil ¹⁴⁵
Coal electricity	Production	1965-2020	BP ¹³⁷
		1960-2020	IEA
Gas electricity	Production	1971-2020	IEA
Traditional biomass	Production	1800-2017	Smil ¹⁴⁵ , OurWorldInData.org
Nuclear electricity	Production	1965-2020	BP ¹³⁷
Hydropower	Production	1890-1964	Smil ¹⁴⁵
		1965-2020	BP ¹³⁷
Biopower	Production	1965-1991	BP ¹³⁷ ("Geo Biomass Other")
		1992-2014	IEA Energy Technology Perspectives 2017 ("Biomass and waste")
		2015-2020	IEA WEOs ("Biomass and waste")
Wind electricity	Production	1984-1989	Wiser et al. ¹⁴⁶ (using capacity factor 0.2, from IRENA 2021)
		1990-2014	IEA Energy Technology Perspectives 2017
		2015-2020	IEA WEOs
Solar PV electricity	Production	1976-1989	Nemet ¹⁴³ (from PCDB)
		1990-2014	IEA Energy Technology Perspectives 2017
Batteries	Production	2015-2020	IEA WEOs
		1996-2016	Ziegler and Trancik ⁸²
PEM electrolyzers	Production	2017-2020	Avicenne (slide presentations online)
		2002-2020	IEA Hydrogen database 2021
P2X fuel	Production	2002-2020	Calculated from electrolyzer data, assuming 50% utilisation.

Table S14: Data sources for Figure 1.

Technology	Conversion efficiency factor
Crude oil	0.25
Coal (primary)	0.6
Gas (primary)	0.6
Electricity	0.9
P2X fuel	0.6
Traditional biomass	0.25

Table S15: Useful energy conversion efficiency factors assumed in Figure 1, based on S.I. 1.8

Figure S17 shows the useful energy cost and production data from Figure 1 plotted as experience curves.

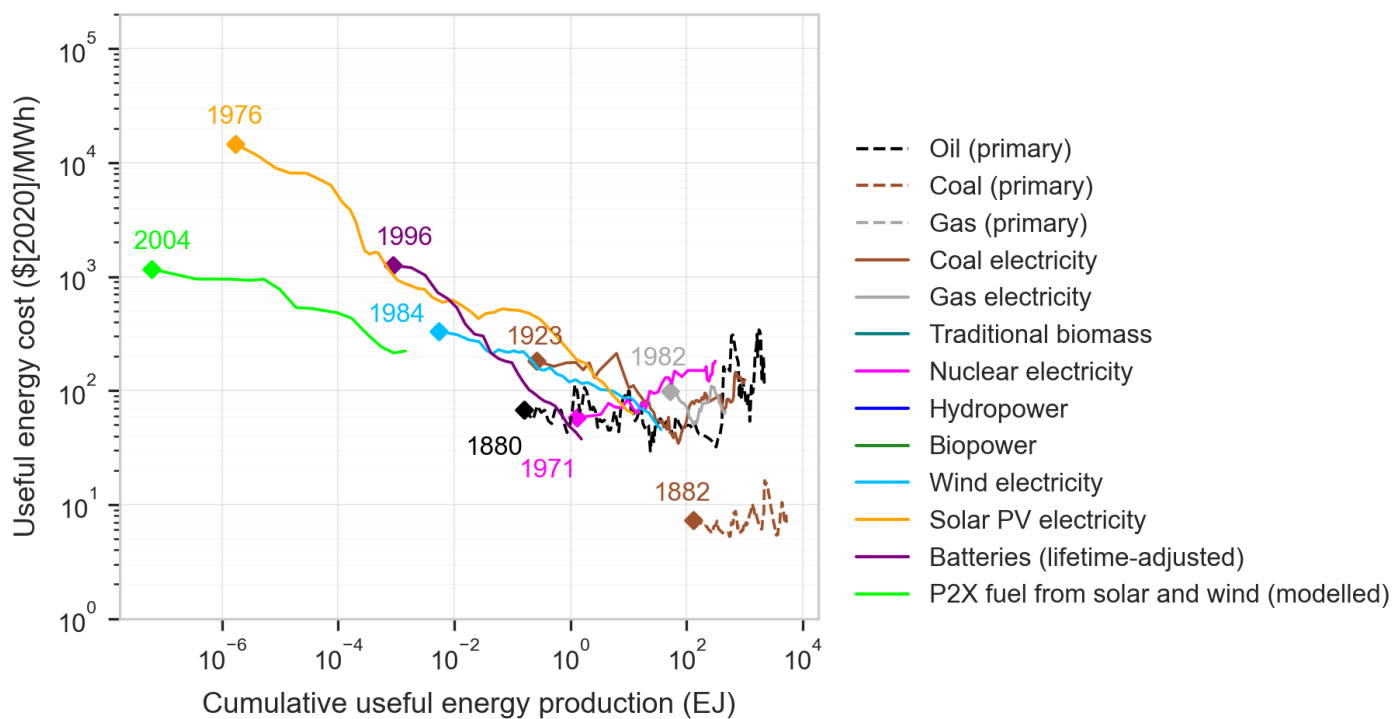


Figure S17: Useful energy costs plotted against cumulative useful energy production for a variety of major energy technologies.

6.2 Oil

To calculate expenditures on oil in different energy scenarios we forecast the price of oil using the AR(1) model. For calibration we use long-run oil price data from BP¹³⁷, plus the 2021 oil price data point from FRED¹⁴⁷ (in order to avoid using an exceptionally low value from during the COVID-19 pandemic as our initial value). There is large variation in the quality, accessibility and hence price of different kinds of crude oil, but we just use the global average value, which in 2021 was 70.86 \$/bbl. One barrel of oil equivalent (BOE) is approximately equal to 6.12 GJ, so the current cost that we use to initialise our scenario forecasting model is $c_0^1 = 11.6$ \$/GJ.

The AR(1) model is fit to data for 1973-2021 (for reasons given in S.I. 5.1) and the estimated parameters are $\hat{\phi} = 0.8128$, $\hat{\sigma}_\epsilon = 0.3037$ and $\hat{\kappa} = 0.4002$. The full data set and the forecast using these estimated parameters are shown in Figure S18.

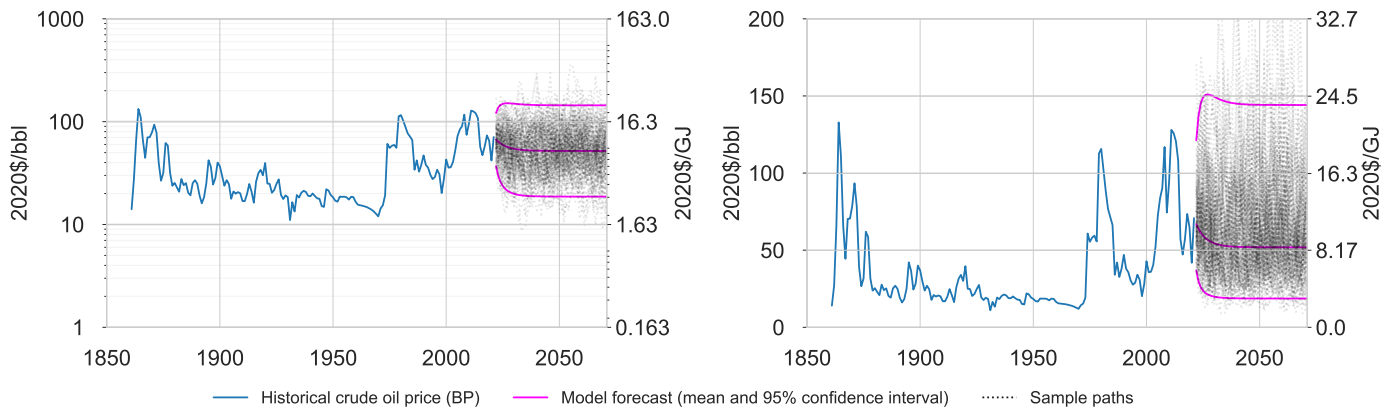


Figure S18: Crude oil price 1861-2021, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

The cost metric we use in the model to calculate expenditures on oil is the price of crude oil, but our energy model considers bulk flows of refined oil. This upgrading process is accounted for endogenously because the energy sector consumes fossil fuels that are used to process raw materials into final energy carriers.

6.3 Coal

Annual coal price data since 1882, collated from various sources, is provided by McNerney et al.¹⁴⁰. The EIA’s website provides price data since 1949 for several different types of coal. The Federal Reserve Bank of St. Louis’s economic data website provides price data since 1980 for Australian coal¹⁴⁷. These are shown in Figure S19.




Figure S19: Coal price data 1882-2020, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

After adjusting for inflation, the price of coal has remained roughly constant for over a century and an AR(1) process provides a good model. Different grades of coal have very different physical properties and are used in different applications¹⁴⁸. The cheapest coal is lignite, used in power generation. Industrial applications require more expensive grades, either bituminous or anthracite. In addition, the price of coal varies depending on location because it is expensive to transport but we ignore this. The energy content of different grades of coal varies significantly: lignite contains around 10-20 MJ/kg, sub-bituminous coal around 19-27 MJ/kg, bituminous coal around 24-35 MJ/kg, and anthracite around 26-33 MJ/kg.

The only use for primary coal in our model is as a direct fuel in industry, so we only need consider the higher grades here (coal-fired electricity is treated separately in S.I. 6.5, and has its own cost model reflecting the low cost lignite fuel). We assume each metric ton of coal contains 30 GJ of energy and calibrate the AR(1) model using the bituminous coal data series. The estimated parameters are: $\hat{\phi} = 0.9499$, $\hat{\sigma}_\epsilon = 0.0902$ and $\hat{\kappa} = 0.0378$. The price of bituminous coal in the US in 2020 was 55.17 \$(2020)/metric ton. We initialize the model with a current price $c_0^2 = 1.84$ \$/GJ (see Table S25).

6.4 Gas

Annual natural gas wellhead price data since 1922 is available from the EIA, and Henry Hub spot price and EU natural gas price data are available from FRED¹⁴⁷. Figure S20 shows the data (1 MMBtu is approximately equal to 1.05506 GJ). Also plotted is the approximate global average price. This is calculated by observing that while the price of gas varies a lot depending on location, around half of the world receives gas at US-level prices, and half at EU-level prices. Therefore, a simple 50-50 mix of these prices approximates the global average price of gas.




Figure S20: Gas price data 1922-2021, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

The real price of gas has remained approximately within an order of magnitude for a century. We use an AR(1) model, which we calibrate on data for 1985-2021 (for reasons given in S.I. 5.1, though we use 1985 as a start date because the structural break associated with the gas shock was slightly delayed relative to oil). The estimated parameters are: $\hat{\phi} = 0.7455$, $\hat{\sigma}_\epsilon = 0.2617$ and $\hat{\kappa} = 0.4028$. In 2021 the average of US and EU prices was 5.8 \$/MMBtu. We therefore take the initial price to be $c_0^3 = 5.5$ \$/GJ (see Table S25).

6.5 Coal electricity

Coal electricity generation data since 1960 is available from the IEA, McNerney et al.¹⁴⁰ and BP. In order to plot around a century of development of the experience curve, we perform a crude estimation of generation from 1923-1960. Annual generation figures in 1960, 1970 and 1980 were around 800, 1500, and 3000 TWh respectively; generation today is over 10000 TWh per year. Although coal power existed for around a century before 1960, generation was very limited compared to today, so we assume a slow taper down from 1960 back to 1930 and negligible generation before that. We estimate pre-1960 cumulative generation to be 12420TWh.

The LCOE of coal-fired electricity varies widely by location, due to large differences in coal transportation costs and environmental regulations and pollution standards. A long time series of LCOE for the US is available from McNerney et al.¹⁴⁰, with Lazard and the EIA providing more recent data also. The time series and experience curve are shown in Figure S21. Underlying coal prices from S.I. 6.3 are also shown on the time series plot, to highlight the connection between the two series since 1960.

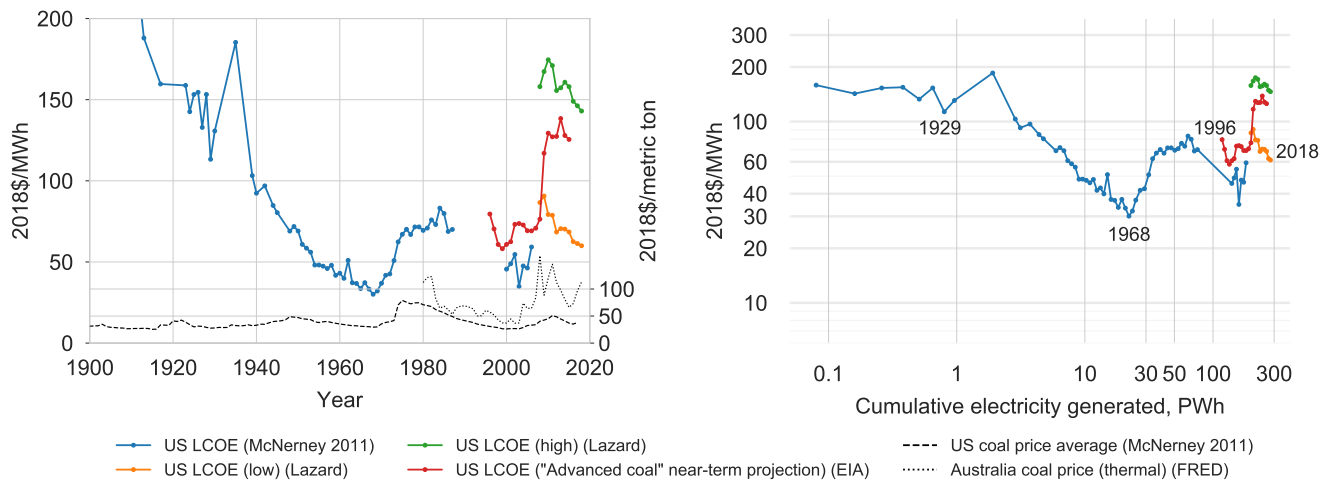


Figure S21: US coal electricity LCOE time series (L) and experience curves (R). Two coal price time series from Figure S19 are also shown on (L). Generation data since 1960 is from the IEA, McNerney et al.¹⁴⁰, BP, and before this it is our own estimate.

As described earlier, the EIA values are near-term projections of what the LCOE of a newly-built plant was expected to be at that time. However, the AEO reports feature several different types of coal technology, which change over time. The data available is so scarce though that we simply use what is available; this is mostly the “Advanced coal” category, but in some years data is only given for the categories “Conventional pulverized coal”, “IGCC” and “Coal”, so we use these.

Despite the ambiguity over precisely which technologies the cost data refers to, it is clear that the LCOE of coal-fired electricity has remained within one order of magnitude for over a century, and even the cheapest plants have remained in the range 30-70 \$(2020)/MWh since 1950. Thus the FDWL model would not be a good fit for this data; an AR(1) process is much more appropriate.

Bearing in mind that model calibration here requires *sequential* data, we fit the AR(1) model to a chained dataset composed of *i*) the largest block of sequential data in the McNerney data (1948-1987) and *ii*) the Lazard low LCOE dataset (2008-2018). The estimated parameters are: $\hat{\phi} = 0.92650$, $\hat{\sigma}_{\eta} = 0.10207$ and $\hat{\kappa} = 0.2056$. Figure S22 shows the associated AR(1) forecast.

We use the Lazard low LCOE data for calibration because, while new-build coal power stations in advanced economies cost around 80-130 \$/MWh, several countries with large coal resources and low pollution standards currently have costs of around 60 \$/MWh. We have not been able to find any long data series for these countries, but BNEF has average LCOE data for 2015-2018. In nominal \$/MWh the ranges in that period are: China: 44-68, India: 48-68, Indonesia: 45-71, Turkey: 45-81, Thailand: 64-83, Japan: 55-80, Vietnam: 68-87. The chained dataset we use for calibration, which ranges from 30-91 \$/MWh, reflects this range quite well. Our cost forecast therefore represents the cheapest, most polluting coal power stations. Since we are focusing purely on minimum possible system costs here, without considering environmental externalities, we use this forecast, with the caveat that it relies on society continuing to accept very high levels of air pollution and carbon emissions; this is a very conservative forecast.*

*To model less polluting coal power stations, more in line with what might be politically acceptable, we would calibrate the




Figure S22: Coal LCOE data 1902-2018, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

model using the EIA data for 1996-2015. This would raise our forecast significantly, as would adding CCS, but we do neither of these.

6.6 Gas electricity

Gas electricity generation data since 1971 is available from the IEA and BP, and in order to produce the experience curve we estimate cumulative generation before then. Annual generation was 696 TWh in 1971, and before this gas did not have much of a history in electricity generation^{149,145}. We estimate cumulative generation before this time to be 7000 TWh.

LCOE data for the US comes from the EIA and Lazard, and some additional global data comes from Colpier and Cornland¹⁴¹. Unfortunately there are no long, consistent time series. Simple open-cycle gas turbines (OCGTs) were developed in the early twentieth century, then more efficient combined-cycle gas turbines (CCGTs) were developed mid-century, and gradually commercialised throughout the 1970s and 80s. OCGTs are simpler, have lower capital costs, efficiency of around 40%, are now used for non-baseload applications (peaking power), and have relatively high LCOE (due to low utilisation). CCGTs are more complex, have higher capital costs, efficiency around 60%, relatively low LCOE, and are used for baseload and load-following power. Our model only considers baseload CCGT plants with high capacity factors.

Fuel accounts for around half of total costs¹⁵⁰, so gas transportation costs influence the LCOE strongly. US costs are generally low because local supply is abundant and transported cheaply by pipelines. Costs in Europe and Japan are higher, reflecting higher transportation costs. We conservatively use the lower US LCOE values in our cost forecasts. The LCOE time series and experience curve are shown in Figure S23. The underlying gas price is also shown.

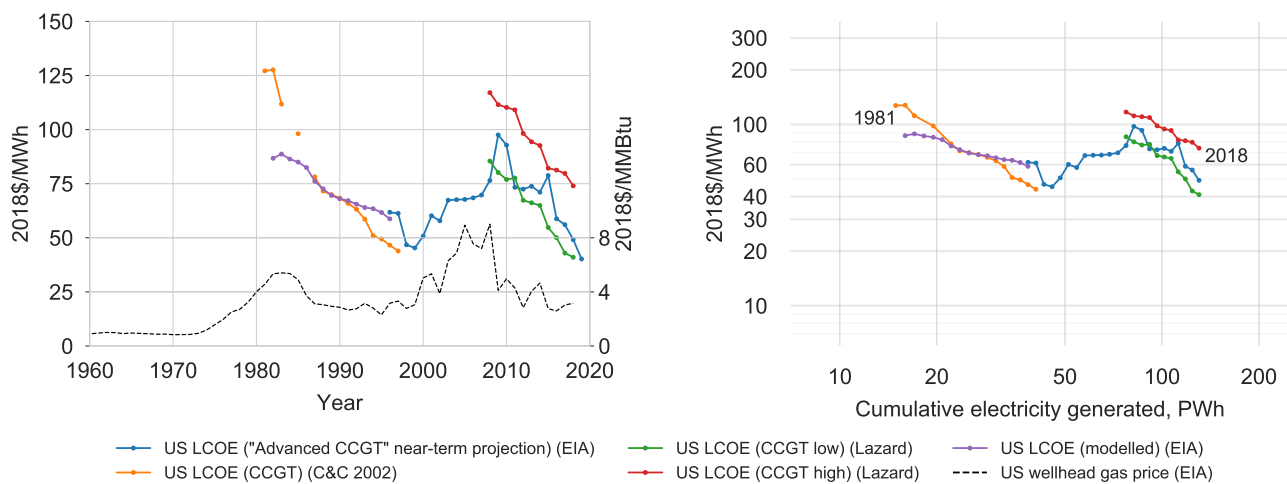


Figure S23: Gas electricity LCOE time series (L) and experience curves (R). A gas price time series from Figure S20 is also shown on (L). Generation data since 1973 is from the IEA and BP.

As with coal power, the EIA's AEO reports feature several different types of technology, which vary over time. We use the "Advanced combined-cycle" category, as this is the most long-lived category in the reports, and the technology itself is relatively cheap and high performance.

We found one other data series from the EIA, which is not very reliable on its own but contains a useful signal when considered alongside other sources. On page 53 of the 1998 AEO edition, Figure 61 shows data from 1982 to 1996 for "Total production costs" of "coal- and gas-fired generating plants". The main difference between this total production cost and an approximate LCOE value is the levelized capital cost component of LCOE. We assume this is constant and match the data to the other EIA series beginning in 1996; this is labelled "US LCOE (modelled) (EIA)" on the figure. We do not use these values for model calibration, just as an indication of the LCOE from 1982-1995.

Due to the general scarcity of data we also include some values from a 2002 study by Colpier & Cornland, despite it not being directly comparable with other data. This paper considered one specific vintage of gas-fired electricity technology, CCGT, from the early stage of commercialisation through to full global deployment, and no further. It captured the trend of high initial costs of this new vintage being reduced over the course of a decade as it scaled up and rapid learning occurred. However, during this expansion phase CCGT only made up a small fraction of the global market, so the high initial costs recorded are not representative of industry-wide average costs. The strong learning trend observed can therefore not be interpreted as a progress trend in the gas-fired electricity industry more widely. As seen when all the data is considered together on the figure, the underlying gas price gives a better indication of the long term trend of gas-fired

electricity costs.

Both the LCOE and the underlying gas price are relatively stable, having remained within one order of magnitude for several decades, so FDWL would not be a good model for this process, and we use the AR(1) model instead. To calibrate it we use the EIA projected LCOE series, as this is the longest series available, and we believe it represents the general trend reasonably well since it tracks the underlying gas price. The AR(1) forecast is shown in Figure S24, and the parameters estimated are: $\hat{\phi} = 0.827042$, $\hat{\sigma}_\eta = 0.130628$ and $\hat{\kappa} = 0.48506$.

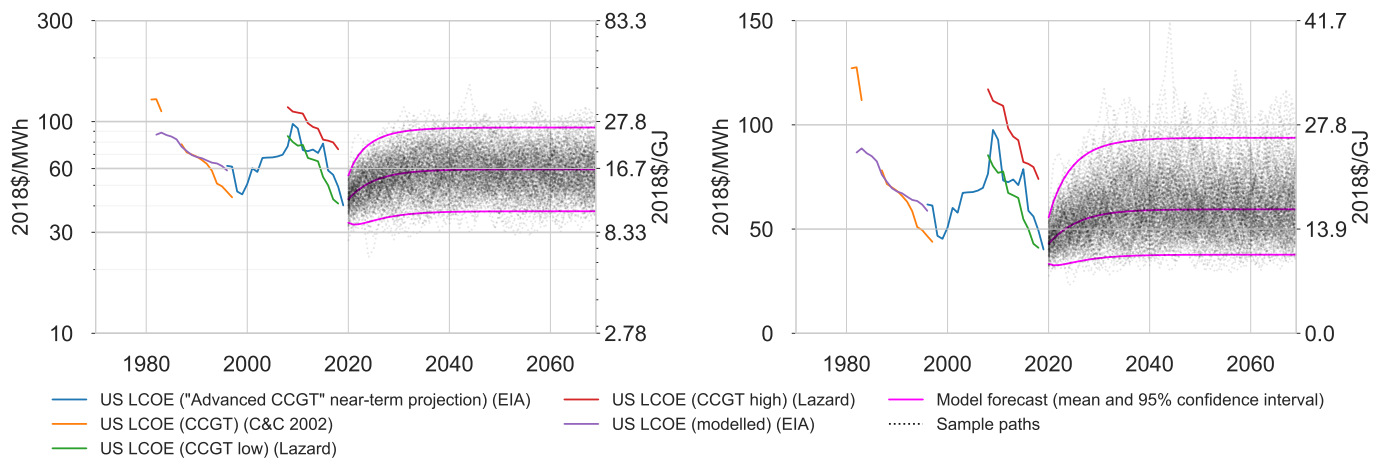


Figure S24: Gas LCOE data 1996-2019. Data: US EIA. Also shown is the AR(1) forecast plus 95% confidence interval, and 200 sample paths.

6.7 Nuclear electricity

Data is available for three different kinds of costs, referring to different technology system scopes. First there is Overnight Construction Cost (OCC) data (also known as overnight capacity cost, or specific investment cost). This is the investment per kW of capacity that would be required to install a nuclear reactor very quickly (i.e. “overnight”) – the cost of instantly buying and installing the required technology. Overnight costs are generally very different from total costs for nuclear power because construction times are long, which leads to large additional interest payments (Interest During Construction). This measure is therefore of limited use in understanding the full costs of nuclear energy. It does however provide some signal regarding underlying technology costs over time, so we include this data for completeness. The main overnight construction cost dataset was compiled by Lovering et al.¹⁵¹, and includes data for demonstration reactors, one-of-a-kind (“turnkey”) reactors, and commercial reactors. There has been a lot of debate about the interpretation and relevance of OCCs, and Koomey et al.¹⁵² and Gilbert et al.¹⁵³ described why they can give a very misleading impression of the total costs of nuclear energy.

Next there is capacity investment cost data; this is the amount of investment per kW of capacity required to install a nuclear reactor in reality, including IDC costs. Interest can be substantial first because plant construction times are long and there are often large delays, leading to cost overruns¹⁵⁴, and second because interest rates reflect wider project risk factors such as local regulation and governance structures. By missing out these elements, OCCs miss out some of the most important cost components in the full nuclear technology system. Grubler¹⁵⁵ analysed capacity investment costs of the French nuclear program and Matsuo and Nei¹⁵⁶ considered the Japanese program.

The final cost data we use is LCOE data, which generally follows a similar trend as capacity investment cost, because nuclear plants have consistently high utilisation. The costs and characteristics of the US nuclear program were analysed by Koomey and Hultman¹⁴², while those in France were analysed by Boccard¹⁵⁷ (based on data from the Court of Audit¹⁵⁸), giving approximate LCOE values over several decades.

We use cumulative generation data from BP as the metric for experience. The time series and experience curve are shown in Figure S25. In addition, Figure S26 shows all the capacity cost data we could find (OCCs and investment costs).


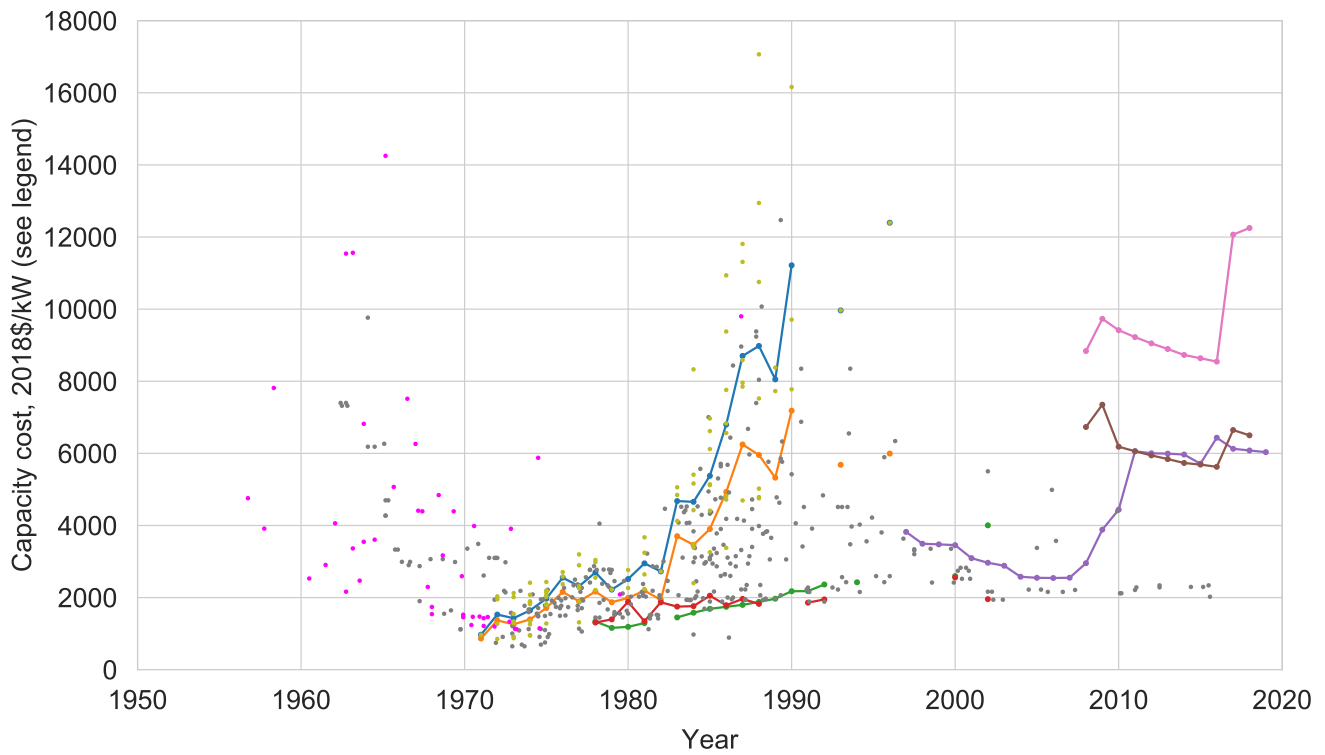


Figure S25: Nuclear power LCOE time series data (L) and experience curves (R).



NB. time refers to start of commercial operations

- US mean capacity investment cost (Koomey & Hultman, 2007), 2018\$/kW
- US mean overnight construction cost (Koomey & Hultman 2007), 2018\$/kW
- France mean overnight construction cost (Grubler, 2010), 2018\$/kW
- France overnight construction cost (Cour des comptes, 2012), 2018\$/kW
- Projected overnight construction cost (EIA AEO cost assumptions), 2018\$/kW
- US capacity investment cost (low) (Lazard), 2018\$/kW
- US capacity investment cost (high) (Lazard), 2018\$/kW
- Overnight construction costs of commercial reactors (global) (Lovering, Yip & Nordhaus, 2016), 2018\$/kW
- Overnight construction costs of demonstrator and turnkey reactors (global) (Lovering, Yip & Nordhaus, 2016), 2018\$/kW
- US capacity investment costs of commercial reactors (Koomey & Hultman, 2007), 2018\$/kW




Figure S26: Nuclear power capacity investment cost data.

Long-term trends in nuclear power are much discussed, but empirically it is clear that on average, globally, costs have increased. The most common understanding of technological learning is that as experience with a technology accumulates, modifications and innovations are added that improve the technology, which lead to cost reductions. However, technological learning can also lead to improvements that result in cost increases. This is the case for nuclear energy, where additional operational experience has resulted in better calibration of risks that were not well understood initially, and consequently better pricing of these risks.

Nuclear energy has half a century of experience in a high skill, technologically advanced, well funded, well supported global industrial sector. It has been continuously active across multiple scales and applications, from marine propulsion to large land based plants. Note that the small, standardised reactors used in marine propulsion, which have been designed, built and operated extensively (more so than land-based power plants), are in many ways similar to “small modular reactors” (SMRs), so it would be surprising if SMRs deviated from the trends displayed in the data here. Despite this vast, varied experience, power generation costs have consistently increased.

Learning occurs in all areas of technology development, from invention to end-use, and for nuclear the most consequential learning has been at the level of the societies in which the technology is embedded. Investors, populations, and politicians have gradually updated their views on the technology – and on the emergent characteristics of the industry that develops the technology – leading to increased safety standards, higher interest rates, and higher costs.

Fitting the FDWL model to the chained series formed from the Koomey and Hultman data and the Lazard low data (see right hand plot in Figure S25) gives parameter estimates $\hat{\omega} = -0.20$, $\hat{\sigma}_{\omega} = 0.15$ and $\hat{\sigma}_{\eta} = 0.14$. Chaining together different combinations of post-commercialisation data (including capacity cost data) gives similar parameter estimates – there is always a negative learning trend. However, there are many highly exceptional features to nuclear power that make its development particularly dependent on policy choices. Therefore, despite the clear long term trend in the data, we attempt to be as optimistic as possible regarding nuclear electricity, on the basis that perhaps a single experience curve encompassing both the 1970s and the 2010s is not the best model. We thus assume that cost escalations have stopped, the current cost is at the very low end of current cost estimates, and the forecast volatility is quite low. We take $c_0^6 = 90$ \$/MWh, $\hat{\omega} = 0.00$, $\hat{\sigma}_{\omega} = 0.01$ and $\hat{\sigma}_{\eta} = 0.02$. Figure S27 shows production in each scenario with the associated cost forecasts given these parameters. We take the technology lifespan to be $L^6 = 40$ years (though varying this choice makes little difference since costs are assumed to be fairly static). Parameter choices are summarised in Table S16.

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.00$	
Experience exponent standard error:	$\hat{\sigma}_{\omega} = 0.01$	
Noise standard deviation:	$\hat{\sigma}_{\eta} = 0.02$	
Current cost:	$c_0^6 = 90$ \$/MWh	$c_0^6 = 70$ \$/MWh
Lifespan:	$L^6 = 40$	

Table S16: Nuclear parameters.



Figure S27: Nuclear electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S25.

6.8 Hydropower electricity

We use global annual hydropower generation data from BP. However, this data begins in 1965, so in order to compute cumulative electricity generation it is necessary to estimate cumulative pre-1965 generation. This is done by using estimates of annual generation available at www.OurWorldInData.org. Generation in 1965 was approximately 1000TWh, and generation in 1940 was approximately 200TWh, so we approximate pre-1965 cumulative generation very crudely by the round number 10000TWh.

There is limited cost data available, in particular no long time series. We use global LCOE data from IRENA¹⁵⁹ and Bloomberg Terminal, but this is only available for the last decade. AEO projected LCOE values are also available for this time, though these refer only to US projects so should be treated with caution, since hydropower is so geographically and politically specific. In addition the AEO reports state that the projected LCOE values assume some seasonal storage capability, and although this technical capability is of course normal, it is not clear if or how the values given have been modified to reflect the value of this service.

Due to the lack of LCOE data we also considered capacity cost data. IRENA has this data since 2010 as usual, and AEO reports give a slightly longer series, from 2005 onwards. The data is shown in Figures S28 and S29.




Figure S28: Hydropower annual generation (L) and capacity cost time series (R).




Figure S29: Hydropower LCOE time series (L) and experience curves (R).

Global LCOEs and capacity costs have been approximately level over the last decade. The US LCOE series shows a steep decline since 2010, but since 2005 capacity costs have actually increased (though most of this is just due to the 2008 financial crash). In this case it seems likely that capacity costs are a better indicator of the general trend, and that the decline in projected LCOEs reported is an artifact of specific modelling assumptions rather than a trend that was ever actually observed.

Note that Wright's law is probably not a very good model for hydropower technology dynamics, because the development of the technology is overwhelmingly dictated by geography – there is relatively little “technology” in this form of generation, compared to other forms. Furthermore, projects are on average very large,

and thus greatly influenced by governmental decision making, so the dynamics of the entire technology system are very different to more granular technologies. As with all large infrastructure projects, significant cost and time overruns are common¹⁶⁰.

Fitting the FDWL model to this very short data series is unsuccessful as this results in experience exponent $\hat{\omega} = 0.60$ and noise standard deviation $\hat{\sigma}_\eta = 0.14$. Due to the extremely large initial cumulative production these values are both very high, and likely completely unrepresentative of the historical development of hydropower costs. Indeed, the plots show that costs are essentially flat, and we have come across no data or anecdotal evidence indicating otherwise. To reflect this trend we take $\hat{\omega} = 0.00$, $\hat{\sigma}_\omega = 0.01$ and $\hat{\sigma}_\eta = 0.01$. Figure S30 shows production in each scenario with the associated cost forecasts given these parameters. We take the technology lifetime to be $L^7 = 100$ years, and use an initial cost value of $c_0^7 = 47$ \$/MWh. Parameter choices are summarised in Table S17.

	Main case
Experience exponent:	$\hat{\omega} = 0.00$
Experience exponent standard error:	$\hat{\sigma}_\omega = 0.01$
Noise standard deviation:	$\hat{\sigma}_\eta = 0.01$
Current cost:	$c_0^7 = 47$ \$/MWh
Lifespan:	$L^7 = 100$

Table S17: Hydropower parameters.




Figure S30: Hydropower electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S29.

6.9 Biopower electricity

For biopower generation data we use the Biofuels and Waste category in the IEA’s Energy Technology Perspectives 2017 report¹⁶¹. The values in this dataset are actual data to 2014 and provisional estimates for 2015-18. However, we treat the latter as true data anyway as it is unlikely to be far from the truth over such a short time horizon. In 1990 generation was around 130 TWh, in 2000 it was around 160 TWh, and in 2010 it was around 370 TWh; we therefore estimate pre-1990 cumulative generation to be 1000 TWh.

Cost data is very limited. We use global LCOE values from IRENA, and US LCOE values from Lazard and the EIA. Due to the general lack of data we also extracted data from Junginger et al.¹⁶², but it is important to note that this is not directly comparable with the other sources. This study tracked the progression of wood chip fuelled CHP plants in Sweden from the very early stages of the industry to nationwide commercialisation. This was a highly directed energy technology program set up specifically to use an abundant, easily managed, local source of biofuel, and the costs are unlikely to be either replicable elsewhere or representative of global trends. The time series and experience curve are shown in Figure S31.




Figure S31: Biopower LCOE time series (L) and experience curves (R).

It is unclear whether an experience curve model or an AR(1) model is more appropriate for this process, since costs do not appear to have dropped significantly in the past three decades. However, use of biopower is so restricted in our model that realistically it doesn’t make much difference which model we use, providing costs remain approximately level. The literature identifies learning rates of 2-15%¹²⁹, so we just pick a value at the very low end of this range and assume small fluctuations around the progress trend. We take $\hat{\omega} = 0.05$, $\hat{\sigma}_{\omega} = 0.01$ and $\hat{\sigma}_{\eta} = 0.02$. Figure S32 shows production in each scenario with the associated cost forecasts given these parameters. We take the technology lifetime to be $L^8 = 30$ years, and use an initial cost value of $c_0^8 = 72$ \$/MWh. Parameter choices are summarised in Table S18.

	Main case
Experience exponent:	$\hat{\omega} = 0.05$
Experience exponent standard error:	$\hat{\sigma}_{\omega} = 0.01$
Noise standard deviation:	$\hat{\sigma}_{\eta} = 0.02$
Current cost:	$c_0^8 = 72$ \$/MWh
Lifespan:	$L^8 = 30$

Table S18: Biopower parameters.




Figure S32: Biopower electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S31.

6.10 Wind electricity

We collected global annual wind electricity generation data from a variety of sources, including BP¹³⁷, IRENA, the IEA, and Wiser et al.¹⁴⁶. We collected cost series for the LCOE of onshore wind, from Wiser et al.¹⁴⁶, Schilling and Esmundo¹⁶³, Lazard¹³⁸, IRENA¹³⁹, and Bloomberg Terminal. Time series, experience curves and data sources are shown in Figure S33.

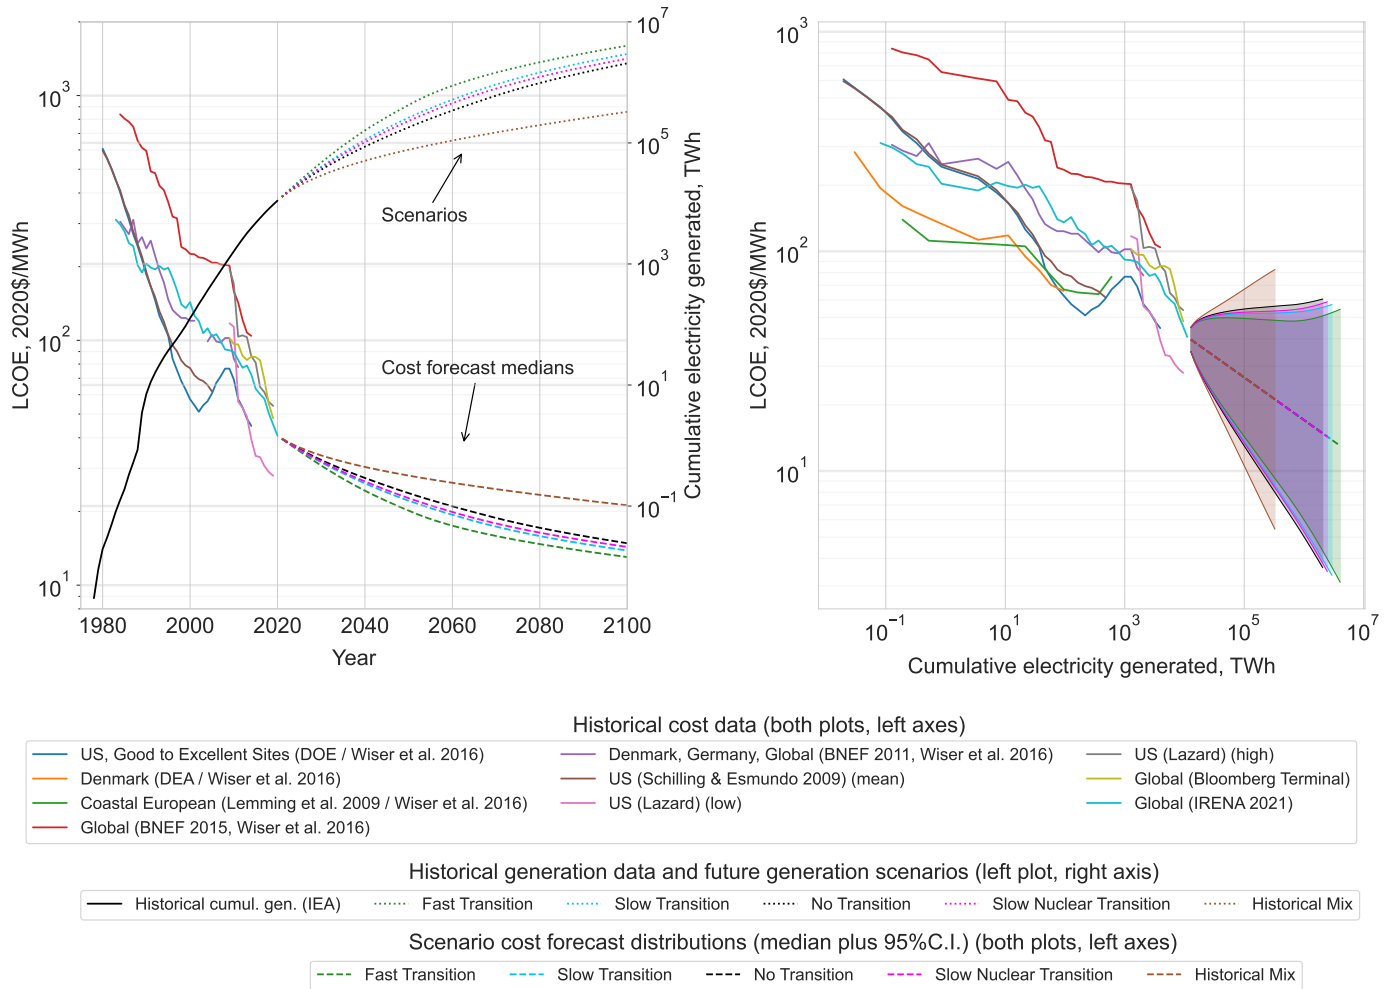


Figure S33: Wind LCOE time series and forecast (L, left axis). Historical generation and future scenarios (R, right axis). Historical experience curves and cost forecasts in each scenario (R).

The final production data series we use is a combination of capacity data from Wiser et al.¹⁴⁶ for 1980-1989 (assuming capacity factor 0.2, taken from IRENA¹³⁹), and IEA generation data for 1990-2020. The final LCOE series we use is from IRENA¹³⁹. These data series give $\hat{\omega} = 0.194$, $\hat{\sigma}_{\omega} = 0.041$ and $\hat{\sigma}_{\eta} = 0.065$. Figure S34 shows production in each scenario with the associated cost forecasts given these parameters. We take the technology lifetime to be $L^9 = 30$ years, and use an initial cost value of $c_0^9 = 41$ \$/MWh. Parameter choices are summarised in Table S19.

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.194$	
Experience exponent standard error:	$\hat{\sigma}_{\omega} = 0.041$	
Noise standard deviation:	$\hat{\sigma}_{\eta} = 0.065$	
Current cost:	$c_0^9 = 41$ \$/MWh	
Lifespan:	$L^9 = 30$	

Table S19: Wind parameters.




Figure S34: Wind electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S33.

6.11 Solar PV electricity

Solar PV generation data for 1976-1989 is from Nemet¹⁴³, for 1990-2014 is from IEA¹⁶¹ and after that from IEA WEO reports. LCOE data is from a wide variety of sources. Historical time series and experience curves are shown in Figure S35, along with generation in each of the five scenarios, plus the implied cost forecast distribution in each case. FDWL is a good model for this technology.




Figure S35: Solar PV LCOE time series and forecast (L, left axis). Historical generation and future scenarios (L, right axis). Historical experience curves and cost forecasts in each scenario (R).

The final production series we use is a combination of Nemet¹⁴³ data for 1976-1989, and IEA data for 1990-2020. The final LCOE series we use is a combination of Nemet¹⁴³ for 1976-2009, and IRENA¹³⁹ for 2010-2020. These data series give $\hat{\omega} = 0.319$, $\hat{\sigma}_{\omega} = 0.043$ and $\hat{\sigma}_{\eta} = 0.111$. Figure S36 shows production in each scenario with the associated cost forecasts given these parameters. We take the technology lifetime to be $L^{10} = 30$ years, and use IRENA's reported 2020 LCOE as the initial cost value: $c_0^{10} = 57$ \$/MWh. For a side case we take a lower value of $c_0^{10} = 49$ \$/MWh, which was reported by BNEF in 2020. Parameter choices are summarised in Table S20.

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.319$	
Experience exponent standard error:	$\hat{\sigma}_{\omega} = 0.043$	
Noise standard deviation:	$\hat{\sigma}_{\eta} = 0.111$	
Current cost:	$c_0^{10} = 57$ \$/MWh	$c_0^{10} = 49$
Lifespan:	$L^{10} = 30$	

Table S20: Solar parameters.




Figure S36: Solar PV electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S35.

Year data published	Publication	Model	PV system floor cost in model, nominal currency	PV system floor cost in model, USD(2020)/kW	Observed global average PV system cost, USD(2020)/kW	Observed 5th percentile PV system cost, USD(2020)/kW
1997	Messner ¹⁶⁴	MESSAGE	1000 USD(1990)/kW	1784	17038	
1999	Barreto and Kypreos ¹⁶⁵	ERIS	500 USD(1998)/kW	757	15113	
2000	Kouvaritakis et al. ¹⁶⁶	POLES	1400 USD(1990)/kW	2498	13465	
2002	de Feber et al. ¹⁶⁷	MARKAL	537 €(1995)/kW	1116	10926	
2004	Anderson and Winne ¹⁶⁸ (from Baker et al. ¹⁶⁹)	E3MG	1000 USD(2004)/kW	1353	9975	
2005	Kypreos ¹⁷⁰	MERGE-ETL	1000 USD(2000)/kW	1465	10131	
2006	Bosetti et al. ¹⁷¹ (from Baker et al. ¹⁶⁹)	WITCH	500 USD(1995)/kW	795	8881	
2007	Bosetti et al. ¹⁷² (from Baker et al. ¹⁶⁹)	WITCH	500 USD(1995)/kW	795	8031	
2009	Bosetti et al. ¹⁷³ (from Baker et al. ¹⁶⁹)	WITCH	500 USD(1995)/kW	795	5112	
2009	Rout et al. ¹⁷⁴	TIMES-G5	2100 €(2000)/kW	3123	5112	
2011	Luderer et al. ¹⁷⁵	ReMIND-R	650 USD(2005)/kW	853	4007	2607
2012	Bibas and Méjean ¹⁷⁶ , Bibas and Méjean ¹⁷⁷	IMACLIM	890 USD(2001)/kW	1274	3021	2085
2014	Criqui et al. ¹⁷⁸	POLES	1100 USD(2005)/kW	1444	2393	1422
2014	Criqui et al. ¹⁷⁸	POLES	2000 USD(2005)/kW	2626	2393	1422
2015	Luderer et al. ¹⁷⁹	REMIND	500 USD(2005)/kW	656	1823	1206
2016	Emmerling et al. ¹⁸⁰	WITCH	400 USD(2005)/kW	525	1657	1032
2017	Muratori et al. ¹⁸¹	GCAM	1680 USD(2010)/kW	1995	1432	837
2017	Creutzig et al. ¹⁸²	REMIND	450 USD(2015)/kW	491	1432	837
2017	Creutzig et al. ¹⁸²	REMIND	225 USD(2015)/kW	245	1432	837
2018	Carrara et al. ¹⁸³	IMAGE	433 USD(2015)/kW	472	1223	806
2018	Carrara et al. ¹⁸³	POLES	619 USD(2015)/kW	675	1223	806
2020	Luderer et al. ¹⁰⁶	REMIND	450 USD(2015)/kW	491	883	572

Table S21: Solar PV installed system floor costs in models, and observed data. Dates in the first two columns may vary because papers are sometimes published online before publication in the print edition of a journal. Deflators and exchange rates used are from <https://fred.stlouisfed.org>. PV system cost data is from annual Trends Reports from the IEA's Photovoltaic Power Systems Programme (<https://iea-pvps.org>) up to 2009, and from IRENA ¹⁸⁴ from 2010 onwards. PV module cost data plotted in Figures 3 and 9a is from the PCDB ¹⁸⁵ (1975-1979) and the Fraunhofer Photovoltaics Report ¹⁸⁶ (1980-2020).

6.12 Batteries

The model has two kinds of batteries available for energy storage, one for daily cycling, the other for a little longer, around three days. We suppose that Li-ion and Va-redox flow batteries fill these roles, though there are other battery chemistries and designs available, and we treat the data presented for these technologies as indicative of general storage cost trends. We deal only in terms of energy storage capacity, GWh. A more comprehensive analysis would also consider charge and discharge rates but that is beyond the scope here.

Li-ion data comes from IEA¹⁸⁷, Schmidt et al.¹⁸⁸, Ziegler and Trancik⁸², plus reports and presentations freely available on the internet from Citibank¹⁸⁹, Pillot¹⁹⁰, BNEF²⁹, Clark¹⁹¹ and Lazard¹⁹². For Va-redox flow batteries, data comes from Schmidt et al.¹⁸⁸ and Lazard¹⁹². Figure S37 shows the time series and experience curves.




Figure S37: Battery cost time series (L) and experience curves (R).

Applying the experience curve concept for electrochemical energy storage is particularly difficult because of the many different chemistries and designs available, and the varied applications – it is not clear where technology boundaries lie. The data shows that since 2010 EV pack costs have converged very quickly to a point in line with the longer term Li-ion trend. We therefore use the longer term trend (“Li-ion (all applications)”) on the plot) and assume that variations in cost for different applications will occur but be roughly in line with this trend. The parameters estimated for Li-ion batteries are $\hat{\omega} = 0.342$, $\hat{\sigma}_{\omega} = 0.073$ and $\hat{\sigma}_{\eta} = 0.116$, so we use these. We take the lifetime of utility-scale grid batteries in the model to be $L^{11} = 12$ years. In addition, we take EV battery lifetime to be 10 years; this is not used directly in the cost analysis, but it is required to calculate the cumulative production of all batteries. We take $c_0^{11} = 310$ \$/kWh as the initial cost value, which is the 2020 value given in the IEA’s WEO 2021. Figure S38 shows production in each scenario with the associated cost forecasts given these parameters.

Further analysis and discussion of grid-scale energy storage technologies may be found in Kittner et al.¹⁹³, while the role of battery technologies in EVs in particular is considered in Nykvist and Nilsson¹⁹⁴, Kittner et al.¹⁹⁵ and Kittner et al.¹⁹⁶. The progress and innovation rates found in these studies are consistent with our modelling.

Based on IEA³⁰ and Pillot¹⁹⁷, we take the installed capacity of EV batteries in 2020 to be 572 GWh, and of grid batteries to be 50 GWh. The cumulative production of all Li-ion batteries in 2020 was 1173 GWh, with the residual made up from batteries in electronics over several decades.

Flow battery technology is much less mature, its future progress more speculative, and the lack of sequential data means we can not calibrate the FDWL model, so we need to estimate some plausible parameters. Schmidt et al.¹⁸⁸ reported a learning rate of $11 \pm 9\%$ when fitting a level experience curve model. The central value of 11% translates to an experience exponent of 0.168 so we use this. This is close to the value for wind electricity (0.194) so we assume that the other two parameters are similar also, and use these: we take $\hat{\omega} = 0.168$, $\hat{\sigma}_{\omega} = 0.041$ and $\hat{\sigma}_{\eta} = 0.065$. This experience exponent is much lower than that observed for Li-ion batteries, but this is plausible because flow batteries are much larger and have few uses beyond large energy system applications, so are less likely to benefit from economies of scale in manufacturing. Lazard’s Levelized Cost of Storage Analysis 2019¹⁹² has the current cost of Va-redox batteries at around 314–550 \$/kWh, so we use an initial cost value of $c_0^{12} = 400$ \$/kWh. We take the technology lifetime to be $L^{12} = 20$ years, and the cu-




Figure S38: Daily battery storage capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S37. The discontinuities between past and future installed capacity in the top row arise because the historical data includes batteries in consumer electronics, which form the majority of all historical battery production, while future installed capacity concerns only batteries used in EVs and power grid applications.

cumulative installed capacity in 2020 to be 0.3 GWh (based on Schmidt et al.¹⁸⁸). Figure S39 shows production in each scenario with the associated cost forecasts given these parameters. Parameter choices are summarised in Table S22.

	Main case (Li-ion)	Main case (flow battery)
Experience exponent:	$\hat{\omega} = 0.421$	$\hat{\omega} = 0.168$
Experience exponent standard error:	$\hat{\sigma}_{\omega} = 0.063$	$\hat{\sigma}_{\omega} = 0.041$
Noise standard deviation:	$\hat{\sigma}_{\eta} = 0.103$	$\hat{\sigma}_{\eta} = 0.065$
Current cost:	$c_0^{11} = 310 \text{ \$/kWh}$	$c_0^{12} = 400 \text{ \$/kWh}$
Lifespan:	$L^{11} = 12$	$L^{12} = 20$

Table S22: Battery parameters.




Figure S39: Multi-day battery storage capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S37.

6.13 Hydrogen and electrolyzers

The model is designed in such a way that scenarios are built in terms of annual quantities of energy from P2X fuels, but costs are computed in terms of the installed capacity of electrolyzers required to produce the P2X fuel. The cost of the electricity used to power the electrolyzers is dealt with separately. We therefore need to know the current installed capacity of electrolyzers, plus their annual P2X fuels output per sector, in order to build scenarios. We also need to know historical electrolyzer cost and production data, to calibrate the FDWL model and forecast costs.

Global hydrogen production is currently around 115 Mt (million tonnes)²⁶. 1 Mt H₂ is 1×10^9 kg H₂, and the lower heating value of hydrogen is approximately 120 MJ/kg (we use the LHV because we assume the latent heat of vaporization of water will not be recovered in most applications). Hence 1 Mt H₂ contains 120×10^9 MJ, or 0.12 EJ, and total current production is around 13.8 EJ. Of this, 60% (8.4 EJ) is pure hydrogen, used mainly in oil refining and ammonia production for fertilisers, and 40% (5.4 EJ) is hydrogen as part of a mixture of gases, used mainly for fuel or feedstock. However, over 95% of hydrogen is produced from fossil fuel sources, the most common method being steam methane reforming (SMR) (around 50%), with coal and oil gasification making up most of the rest. Producing hydrogen by electrolysis of water currently makes up just 2-4% of the total. Despite the current dominance of fossil fuel based methods, decarbonising these would require CCS, so we do not consider them in our model, we only consider electrolysis.

Electrolyzers use electricity to split water and produce hydrogen. This is the reverse of the process that occurs in a hydrogen fuel cell, and the two technologies are closely linked. Indeed some current electrolyzers can be run in reverse, providing fuel cell functionality. In our high renewables scenarios we assume that either some electrolyzers will have this capability and will be used for electric grid backup in future (running in one direction when excess renewable electricity is available and in the other when there is a scarcity); or, if this functionality does not become widespread, that hydrogen can be combusted in turbines at very low cost (since there will be a large stock of fully amortized CCGTs available after gas combustion has been phased out, see S.I. 3.3). Most electrolyzers, however, will be dedicated to hydrogen production for other purposes (transport, industry, and buildings) and will not require this functionality. There are three established types of electrolysis and electrolyzer, with different characteristics and different potential for future development: alkaline electrolysis (AEL), polymer electrolyte membrane electrolysis (PEM) and solid oxide electrolysis cells (SOEC) (see e.g. Buttler and Spliethoff¹⁹⁸, Carmo et al.¹⁹⁹, IRENA²⁰⁰ for comprehensive reviews).

AEL is a mature technology that has been used commercially since the 1920s, especially in chlorine and fertiliser production. PEM and SOEC were invented in the 1960s and 70s respectively. Both are currently more expensive than AEL and have been used far less. Before SMR and gasification methods became widespread in the 1950s, AEL was the primary method of hydrogen production. It still accounts for almost all hydrogen produced by electrolysis today.

Around 20 GW of cumulative AEL capacity has been installed, though much of this has now been decommissioned^{201,188}. Early installations were placed near hydropower facilities to take advantage of cheap, reliable electricity, but this practice stopped when the cheaper fossil-based methods took over. Both PEM and SOEC have much lower cumulative installed capacities, probably around 80 MW and 0.3 MW respectively. These are figures given by the IEA for energy applications only, and it is not clear how widely the technologies have been deployed in other areas, but due to their higher costs relative to AEL, it is believed that they have only occupied very small niches with specific performance requirements, for example in science, R&D and hi-tech environments.

One of the main drawbacks of AEL, especially in the context of operating with highly variable renewable energy sources as input, is its relative lack of flexibility. It currently takes 1-5 minutes to achieve full load from an already warm, pressurized standby state, while PEM takes just seconds: less than 10 seconds for MW-scale installations. This lack of flexibility could be addressed either by advances in the technology itself (which reports indicate are underway) or by coupling it with a battery storage system (the costs of which are falling rapidly also). PEM costs have dropped a lot in the last decade as the technology has begun to reach commercial scale. SOEC is still mostly at the R&D stage, though it is attracting a lot of interest due to its high conversion efficiency and ability to run in reverse as a solid oxide fuel cell (SOFC), generating electricity from hydrogen. PEM fuel cells also exist, making reversible PEM electrolyzers a possibility in future.

Due to the complicated historical development of this group of technologies, and their different technical characteristics, there is no obviously best way to apply the experience curve concept. AEL is orders of magnitude more mature than PEM and SOEC, so treating them all as a single technology and using the combined cumulative experience to construct an experience curve would mean that any trends that may exist for PEM and SOEC alone would be suppressed by the AEL trend (i.e. any experience-cost trend would be obscured

by the much larger AEL experience). Furthermore, knowledge spillovers, at least between AEL and the other two technologies, are likely to be small, since their different characteristics mean that they perform different functions and are part of different technology clusters – they are not usually substitutes. Hence we believe it is reasonable to treat AEL and PEM/SOEC as separate technologies, with separate experience curves.

We have collected as much data as possible on costs and cumulative experience, but it is relatively patchy overall, and does not lead to definitive conclusions about how best to model future costs. Before considering the best strategy it is useful to consider the data. There is not enough data on SOEC to make meaningful plots so we just consider AEL and PEM.

For production data we referred to: Schmidt et al.¹⁸⁸, Glenk et al.¹⁴⁴, IEA²⁶, and IEA⁸⁹, and the accompanying datasets, plus Schoots et al.²⁰¹, Buttler and Spliethoff¹⁹⁸ and Staffell et al.⁹⁶. Of these, Schmidt et al.¹⁸⁸ and Schoots et al.²⁰¹ consider the full history of AEL applications, while all the others consider electrolyzer capacity for energy purposes only, i.e. power-to-gas (PtG) systems – this is an important difference.

For cost data we referred to Glenk et al.¹⁴⁴, Glenk and Reichelstein²⁰² and Schmidt et al.¹⁸⁸, and the accompanying datasets, plus IEA¹², IEA¹⁰², and IRENA²⁰³. In all cases costs refer to electrolyzer system costs. A full PtG system costs about 15% more²⁰², which is typically well within the range of cost estimates anyway, so the difference is not important. We plot the data in Figures S40-S42. Our final cost data series for PEM electrolyzers is the combination of the Glenk et al.¹⁴⁴ data, which runs from 2004-2019, plus one final 2020 cost value taken from IEA¹⁰².

Figure S40 shows the cost time series and experience curve for AEL since the 1950s. Despite the much better data coverage for PtG applications since 2003, the experience curve is not particularly meaningful because recent PtG capacity additions form such a small part of the historical total, and show a great deal of variation. IEA²⁶ lists 2018 capital costs for a full AEL PtG plant as being in the range 500-1400\$/kW. However, recent quotes reported by BNEF indicate that costs may be as low as 200\$/kW already in some places²⁰⁴.

If instead we suppose that pre-1990s learning has played only a minimal role in AEL's subsequent progress in PtG applications (which is plausible), then we see a slightly different picture because cumulative experience in these applications is much lower.* Since we assume PEM systems have no other experience besides PtG systems, this perspective gives a fairer comparison between the two technologies. Figures S41 and S42 show the PtG plots for AEL and PEM systems.

We observe that progress in PEM PtG has on average been faster than in AEL PtG, and the experience curve appears to be steeper and more consistent. However, it is clear that in both cases the data itself is very volatile. This is for two reasons, both a consequence of the lack of maturity of the technologies. First, there are relatively few cost observations, so that any calculated average cost is very volatile. Second, production additions are very lumpy, sometimes arriving in very small increments, sometimes much larger. Calibrating the FDWL on such volatile data results in forecasts with very large uncertainty.

Attempting to calibrate the FDWL model using the AEL PtG mean cost data does not give good results because the fluctuations are so large and there is so little data that the model is unable to pick out a progress trend: $\hat{\omega} = -0.045$, $\hat{\sigma}_\eta = 0.439$ and $\hat{\sigma}_\eta = 0.368$. Calibrating the model using the PEM PtG cost data from Glenk et al.¹⁴⁴ is more successful, resulting in an estimated experience exponent of $\hat{\omega} = 0.129$ with standard error $\hat{\sigma}_\eta = 0.067$, and noise standard deviation $\hat{\sigma}_\eta = 0.201$, so although the uncertainty is still large, there is also a strong progress trend.

Given the wide range of costs, cumulative production and characteristics of the different technologies, we believe that the most justifiable strategy for producing data-driven, quantitative distributional cost forecasts is to use the PEM PtG dataset, with the understanding that there is a good chance this trend may also encompass SOEC technologies as they progress. Currently though, PEM has the high flexibility required to absorb variable electricity sources, it is safe and reliable, and has good lifetime and efficiency characteristics, so it is a very promising candidate. Note that the lowest current cost estimate of AEL is already 80% cheaper than the initial PEM cost estimate we use, so we are again being highly conservative in our cost estimates.

*In fact it seems likely that pre-1990s progress in all three technologies may have played a similarly small role in their overall progress in PtG systems. After all, they were all invented many decades earlier, and had years of experience at the R&D stage. Also, the main commercial application of early AEL systems involved continuous, very high load factor operation, which is quite different from the kind of flexible operation being considered for modern PtG applications.




Figure S40: AEL electrolyzer cost time series (L) and experience curve (R).




Figure S41: AEL PtG/PtL electrolyzer cost time series (L) and experience curve (R).




Figure S42: PEM PtG/PtL electrolyzer cost time series (L) and experience curve (R). “GMR 2021” means Glenk et al.¹⁴⁴, and “G&R 2019” means Glenk and Reichelstein²⁰². (Note that the 2020 PEM cost value is taken from IEA¹⁰², because the Glenk et al.¹⁴⁴ data only runs to 2019.)

Now that we have chosen to use PEM as our electrolyzer technology, we must consider the initial capacity that exists and the initial hydrogen production per sector. The current installed capacity is around 80.24 MW for PEM in energy applications. To get this figure we used the data in the IEA⁸⁹ database, but excluded one very large data point (8.931MW) from 2000, as it is completely out of character with subsequent data (it may be measuring something other than PEM PtG systems). In lieu of this single point, we added two other data points we found that are much more in line with subsequent data. The first is for 1992 (0.002MW), the second is for 2000 (0.001MW), and they were found in Buttler and Spliethoff¹⁹⁸ and IEA²⁶, respectively.

The total installed PEM capacity, 80.24 MW, is approximately equal to the total cumulative production of these electrolyzers, because so few have been decommissioned. 80.24 MW of capacity is equivalent to 288.9 GJ/hour. This is the maximum rate of electricity consumption. In this model we assume electrolyzers are operated at 50% utilisation rate on average, so the annual electricity consumption of PEM electrolyzers is around 0.001265 EJ/year. This is the electrical input though, so assuming a conversion efficiency of 70%, this gives around 0.0008857 EJ/year of hydrogen.[†] This is such a tiny amount that the precise value is not important, just the approximate order of magnitude. We split this electrolytic hydrogen evenly between the three end-use sectors to initialise the quantity of P2X fuel consumed by each, from which to begin building out scenarios.

We take $\hat{\omega} = 0.129$, $\hat{\sigma}_{\omega} = 0.067$ and $\hat{\sigma}_{\eta} = 0.201$. Figure S43 shows production in each scenario with the associated cost forecasts given these parameters. PEM electrolyzer stack lifetimes (i.e. operating hours) are currently in the range 30000-90000 hours, and are expected to rise by 2030 to 60000-90000 hours²⁶. We assume a lifetime of 70000 hours, which, at 50% utilisation rate, translates to around 16 years. We therefore take the technology lifetime to be $L^{13} = 16$ years, and use an initial cost value of $c_0^{13} = 1313$ \$/kW, which is the value for 2020 given in the IEA's WEO 2020. Parameter choices are summarised in Table S23.




Figure S43: Electrolyzer installed capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figures S41 and S42.

[†]Conversion efficiencies of current PEM electrolyzers are around 60%. However, the long term efficiency of PEM is expected to be around 70%²⁶. We use a single, fixed value of 70% throughout the model, because it is only in the long term that electrolysis produces significant quantities of hydrogen. In addition, utilisation rate and efficiency play equivalent roles in the hydrogen production calculation, so it would be easy to compensate for the slightly high initial efficiency value if we wanted, by assuming slightly higher initial utilisation. This is unnecessary though, and we prefer simply to use single parameter values that remain constant for the entire model horizon. In addition, SOEC efficiency is currently around 80% and is expected to increase to around 90%, so there is plenty of opportunity for improvement.

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.129$	
Experience exponent standard error:	$\hat{\sigma}_{\omega} = 0.067$	
Noise standard deviation:	$\hat{\sigma}_{\eta} = 0.201$	$\hat{\sigma}_{\eta} = 0.100$
Current cost:	$c_0^{13} = 1313 \text{ \$/kW}$	
Lifespan:	$L^{13} = 16$	

Table S23: Electrolyzer parameters.

6.14 Electricity networks

Power grid costs are dealt with differently than other technologies. Due to limited data and the relatively small contribution to total system costs, we do not use a stochastic cost forecasting model, but simple deterministic costs instead, that scale with total electricity generation. In addition to the investments shown here, the model also includes large expenditures on batteries, electrolyzers and overcapacity, to deal with the intermittency of VRE sources, as described in detail in S.I. 3.7.

Figure S44(L) shows total electricity generation historically, and in each of our scenarios. Observed data here is from IEA¹⁶¹ for 2000-2014, and thereafter from annual IEA WEO reports.

Figure S44(R) shows annual investments in electricity networks historically, plus our modelled investments in each scenario. Observed data is from IEA¹⁶¹ for 2000-2009; from the IEA's 2019 World Energy Investment report²⁰⁵ for 2010-2018; and from IEA¹¹⁰ for 2019-2020.




Figure S44: Total electricity generation (L) and annual investments in electricity networks (R) in each scenario.

6.15 Technology data summary

Table S24 summarises all the technology data used in the model, and Table S25 summarises the unit conversions.

Technology	Cost, c_0 (2020\$)		Cumulative production, z_0		Annual production, E_0		Lifespan, years	Cost model parameters			
	Std. units	\$/GJ	Std. units	EJ	Std. units	EJ/yr		Mean-reversion $\hat{\phi}$	Constant $\hat{\kappa}$	Noise $\hat{\sigma}_\epsilon$	Long-term mean, $e^{\hat{\mu}}$
AR(1) technologies											
Oil ¹	70.86 \$/bbl	11.6	-	-	3229 Mtoe	135.2	-	0.8128	0.4002	0.3037	8.476
Coal ²	55.17 \$/metric ton	1.84	-	-	910 Mtoe	38.1	-	0.9499	0.0378	0.0902	2.126
Gas ³	5.8 \$/MMBtu	5.5	-	-	1421 Mtoe	59.5	-	0.7455	0.4028	0.2617	4.8690
Coal electricity	60 \$/MWh	16.7	310624 TWh	1118	9911 TWh	35.7	-	0.927	0.206	0.102	16.4
Gas electricity	50 \$/MWh	13.9	142762 TWh	514	6356 TWh	22.9	-	0.827	0.485	0.131	16.5
Wright's law technologies											
Nuclear	90 \$/MWh	25	95782 TWh	345	2790 TWh	10.0	40	0.000	0.010	0.020	0
Hydropower	47 \$/MWh	13	144187 TWh	519	4236 TWh	15.2	100	0.000	0.010	0.010	0
Biopower	72 \$/MWh	20.0	10350 TWh	37.3	709 TWh	2.55	30	0.050	0.010	0.020	3.4%
Wind	41 \$/MWh	11.3	11158 TWh	40	1596 TWh	5.75	30	0.194	0.041	0.065	12.6%
Solar PV	57 \$/MWh	15.7	3683 TWh	13.3	833 TWh	3.0	30	0.319	0.043	0.111	19.8%
Daily batteries	310 \$/kWh	86000	1173 GWh	0.00422	-	-	12	0.421	0.063	0.103	25.3%
Multi-day storage	400 \$/kWh	111100	0.3 GWh	10.8×10^{-7}	-	-	20	0.168	0.041	0.065	11.0%
Electrolyzers	1313 \$/kW	364722	80.24 MW	2.89×10^{-7} EJ/hr	-	-	10	0.129	0.067	0.201	8.6%
Other technologies											
Electricity networks	10.4 bn\$/PWh										

Table S24: Technology parameters used in main case analysis. All values refer to the model base year, 2020.

¹ Direct-use oil in transport, industry and buildings sectors only

² Direct-use coal in industry and buildings sectors only

³ Direct-use gas in industry and buildings sectors only

Technology	Relevant units and conversions
Crude oil	<p>bbl: barrel of oil (unit of volume) boe: barrels of oil equivalent (unit of energy) Mtoe: million tonnes of oil equivalent (unit of energy) 1 boe = 6.12 GJ 1 Mtoe = 0.041868 EJ $\therefore 1 \text{ \\$/boe} = \frac{1}{6.12} \text{ \\$/GJ} = 0.163 \text{ \\$/GJ}$ Average price in 2021 ¹⁴⁷: 70.86 \\$/bbl This is equivalent to approximately 11.6 \\$/GJ = 11.6 bn\\$/EJ</p>
Bituminous coal (i.e. non-power-sector coal)	<p>1 short ton = 0.907185 metric tons = 907.185 kg (unit of mass) Specific energy: 24-35 MJ/kg We use 30 MJ/kg = 30 GJ/metric ton Bituminous price in 2018, US ²⁰⁶: 59.43 \\$/short ton = 65.51 \\$/metric ton Anthracite price in 2018, US ²⁰⁶: 99.97 \\$/short ton = 110.2 \\$/metric ton Price in 2018, Australia ¹⁴⁷: 113.23 \\$/metric ton Bituminous US value is equivalent to 2.18 \\$/GJ</p>
Natural gas	<p>MMBtu: million Btu (unit of energy) 1 MMBtu = 1.05505585 GJ (IEA) Price in Jan 2019, US (FRED): 3.08 \\$/MMBtu Price in Jan 2019, Europe (Statista): 6 \\$/MMBtu Price in 2019, Japan (Statista): 7.4 \\$/MMBtu Lowest value is equivalent to approximately 3 \\$/GJ = 3 bn\\$/EJ</p>
Electricity	<p>1 TWh = 0.0036 EJ 1 MWh = 3.6 GJ 1 kWh = 3.6 MJ 277.8 TWh = 1 EJ $\therefore 1 \text{ \\$/MWh} = \frac{1}{3.6} \text{ \\$/GJ} = 0.2778 \text{ \\$/GJ}$ and $1 \text{ \\$/kWh} = \frac{1}{3.6} \text{ \\$/MJ} = 0.2778 \text{ \\$/MJ} = 277.8 \text{ \\$/GJ}$ E.g. LCOE: 50 \\$/MWh = 13.90 \\$/GJ = 13.90 bn\\$/EJ E.g. EV battery packs: 137 \\$/kWh \approx 38000 \\$/GJ E.g. Li-ion grid storage: 310 \\$/kWh \approx 86000 \\$/GJ E.g. VRF grid storage: 400 \\$/kWh \approx 111000 \\$/GJ</p>
Electrolyzers	<p>Current cost is around 1313 \\$/kW 1 kW = 1 kWh/h = 10^{-9} TWh/h = 3.6×10^{-12} EJ/h $\therefore 1313 \text{ \\$/kW} = 1313 \times 10^{12} / 3.6 \text{ \\$/EJ/h} \approx 364722 \text{ \\$/GJ/h}$</p>

Table S25: Summary of cost and unit conversions.

7 Results

7.1 Main case results

Calculating the total costs of each scenario over one instance of the next T years requires generating the complete ensemble of stochastic variables in the model: T noise draws for each AR(1) technology, plus one experience exponent draw and $T+1$ noise draws for each FDWL technology. For $T = 50$ the resulting ensemble consists of 666 values.

We generated this complete ensemble $M = 100000$ times, with the parameters shown in Table S24 (this takes around 50 minutes on a 2015 MacBook Pro). For each ensemble we computed the implied cost sample path for each technology (shown in S.I. 6, e.g. Figures S19 and S27). We then calculated the annual system cost sample at each time for each scenario (Eq. 49), thus obtaining the annual cost distributions, $W_t^{scenario}$. Figure S45 shows the evolution of these distributions for the five scenarios, and Figure S46 shows the median annual spend on each technology in each scenario.*




Figure S45: Annual system cost distributions, $W_t^{scenario}$, for each scenario. Medians are shown as solid lines, expectations as dashed lines, and the shaded regions represent the 95% confidence intervals.

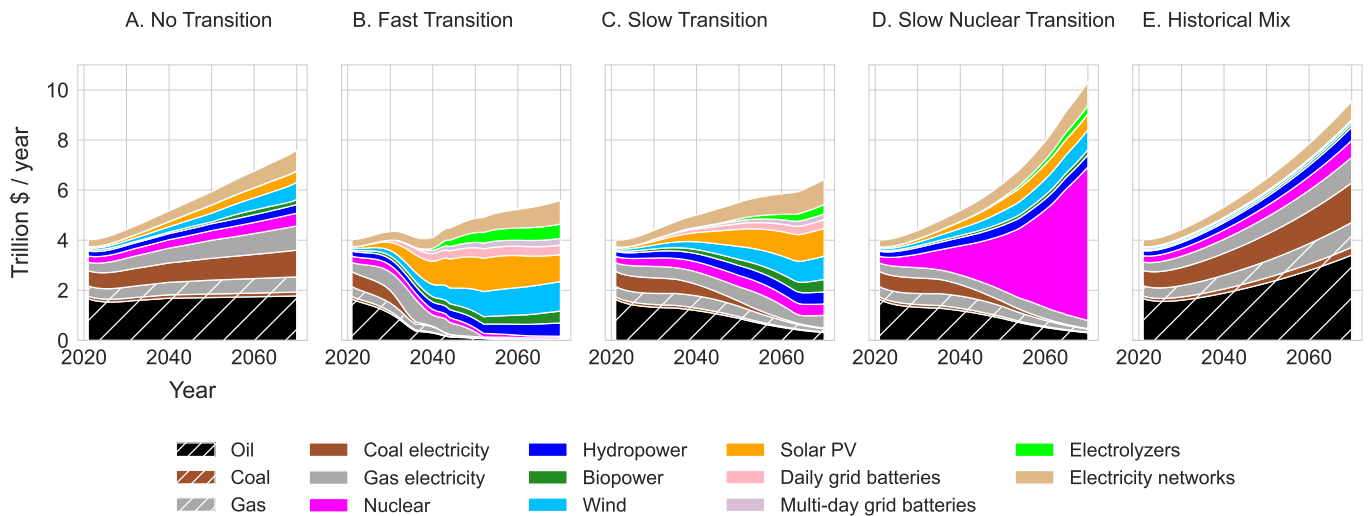


Figure S46: Median annual spend on each technology in each scenario.

*Note that the medians in Figure S45 differ from the sum of the medians in Figure S46 because the sum of medians of distributions does not in general equal the median of a sum of distributions.

In the Fast Transition scenario the annual system costs drop sharply and with high probability in the early 2030s as low cost renewables displace fossil fuels throughout the power system and light-duty vehicles. Annual costs rise again in the late 2030s as large quantities of green fuels are required to displace fossil fuels in industry and hard-to-decarbonise applications, but both mean and median annual costs are still lower than the other scenarios. After 2050, due to the way parameter uncertainty affects the mean and median of cost forecast differently in the FDWL model (as noted in S.I. 5.2), these diverge strongly.

We then picked a discount rate r and, for each of the M ensembles, computed the present discounted system cost sample for each scenario (Eq. 50), thus obtaining the distributions $V^{scenario}(r)$. We repeated this for several different discount rates to observe how they vary with r . These are shown in Figure S47 (L) for a few r values.

For each of the M ensembles, we subtracted the present discounted cost sample of each scenario from the present discounted cost sample of the No Transition scenario. This gives the relative NPC distribution of each scenario for discount rate r , from which we then compute the expectation and median. By repeating this for several discount rates we observe how the expected relative NPC of each scenario varies with r . This is shown in Figure S47 (R). The corresponding analysis for medians is presented in Figure S48.




Figure S47: Present discounted cost distribution of each scenario at varying discount rates, with expected values shown as dashed lines (L). Expected net present cost of transition for each scenario relative to No Transition (R).

Finally, for each fixed discount rate r , we pairwise sampled the present system cost distributions $V^{scenario}(r)$ a large number of times (100000 times), to obtain approximate values for the probability that one scenario is cheaper than each of the others. Results are shown in Figure S49.

The expected relative net present cost of a rapid green energy transition is negative at all discount rates up to around 25%, and lower than any other scenario modelled here. Therefore, in terms of energy system and technology costs alone, implementing the Fast Transition scenario is a bet worth taking. The expected payoff of the bet (i.e. the expected NPC savings up to 2070) is around 5-15 trillion dollars. In addition, Figure S49 shows that the probability that this bet will pay off at all (i.e. that the NPC of the Fast Transition will be lower than that of another scenario, in a pairwise comparison) is around 80%, when compared against any other scenario. This is by far the safest bet, out of all the scenarios. To put it another way, implementing a slower energy transition, as in the Slow Transition scenario, or continuing with the current fossil-fuel-based energy system, as in the No Transition scenario, both have around an 80% chance of being *more expensive* than rapidly transitioning to a green energy system based on the technologies laid out in the Fast Transition scenario. Furthermore, this analysis does not include climate damages, so the full economic costs of No Transition, Slow Transition, Slow Nuclear Transition, and Historical Mix would be much larger than the system costs given here. A fuller analysis would increase their costs dramatically, either due to climate damages resulting from increased emissions, or due to enormous mitigation costs on top of those considered here (CCS for example). See S.I. 8.8.1 for an estimate of these additional costs.

It is important to understand that while these results may initially appear to be an obvious consequence of the cost forecast models used, they are not, and there is no way the results could have been known *a priori*.




Figure S48: Present discounted cost distribution of each scenario at varying discount rates, with median values shown as dashed lines (L). Median net present cost of transition for each scenario relative to No Transition (R).

To address the notion that these results are “simply a function of the cost models chosen”, we first note this is in some sense trivially true, since all model results are a function of model structure and the assumptions chosen. Yet this criticism is not levelled at more complex models, precisely because their complexity obscures the dependency between model outputs and inputs. There are cases where this dependency becomes obvious though. For example, as discussed in S.I. 8.5, IAM results are sensitive to the cost model used for PV (and other technologies), in particular their assumed floor costs (the lower the PV floor cost, the more PV is built). Yet their results are still considered valid and informative. So the more relevant question is whether the cost models used are the most appropriate ones. In our case, due to the statistical testing these models have undergone, we believe they are.

More importantly though, while the reason that IAM results can not be known *a priori* is their sheer complexity, the reason our results can not be known *a priori* is more subtle. Our cost models are calibrated on observed technology data, so our technology cost forecast distributions, and the way in which they sum to form system cost distributions at different years in the future, depend directly on this data. Different observed data would yield different results. In particular, the relationship between the historical progress trend and the historical volatility, is not at all obvious, and it is these parameters that combine to generate forecast uncertainty, which is especially important in the expected NPC calculation. Of particular interest is the relatively low uncertainty in the annual system cost forecast in the Fast Transition scenario in the mid-2030s, shown in Figure S45. Without this relatively low uncertainty early on, it is quite possible that the relatively high uncertainty in later years would have caused the NPC results to be rather different.

7.2 Sensitivity to time horizon T

How do the results depend upon the time horizon of the model? Recall, scenarios are constructed up until 2100, but the NPC calculations may be performed on any time horizon up to that point. On the one hand, the NPC calculation should allow enough time for the benefits of falling technology costs to be felt, i.e. enough time for “building solar on solar” – powering a doubling of the global economy by 2050 with clean energy, and using this growth to drive costs down further. On the other hand, forecast uncertainty increases with time, so looking too far ahead is futile. The latter point is captured by the difference in the mean and median cost forecasts given by the FDWL model. The median follows the deterministic experience curve trend, while the mean initially follows this quite closely but then, due to the log-normal cost forecast structure, diverges towards infinity. The mean and median of the annual system cost distribution also begin to diverge after thirty years. To explore the effects of these issues on the NPC results, we repeated the calculations for time horizons from 2050-2090. The results are shown in Figure S50.

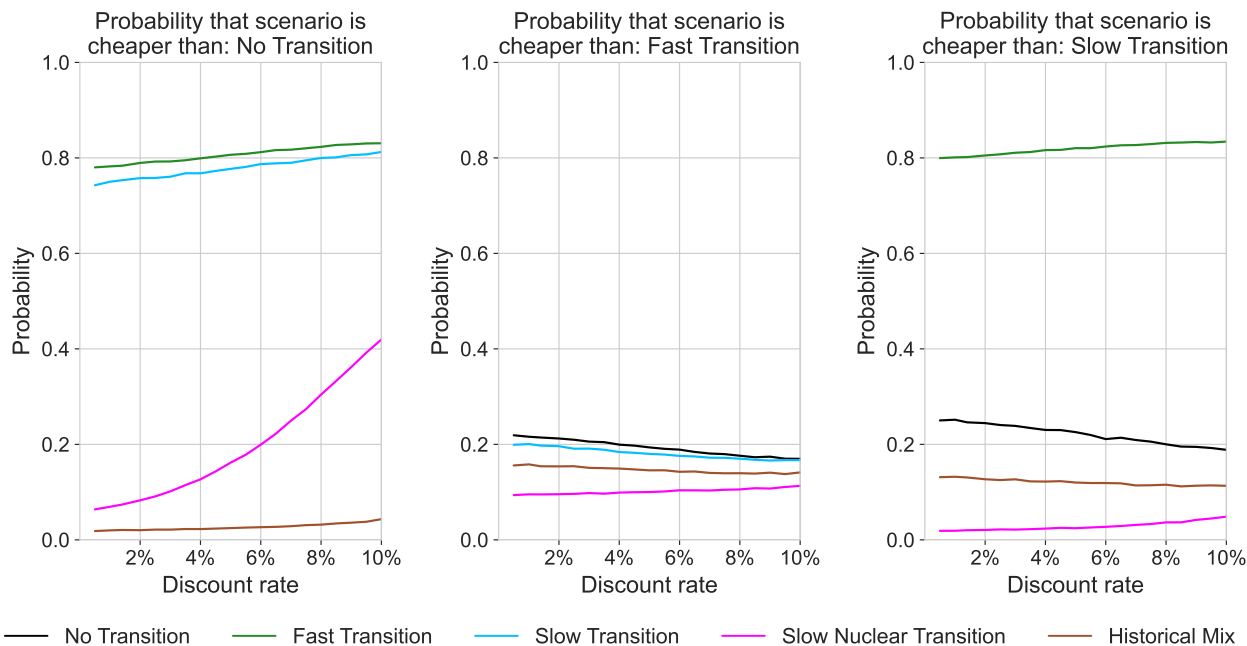


Figure S49: Probabilities that scenario NPCs are lower than those of No Transition (left), Fast Transition (middle), and Slow Transition (right), for varying discount rates.




Figure S50: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the five scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090.

7.3 Sensitivity to system growth rate

The central assumption made in this work is that the global supply of useful energy services in end-use sectors grows exogenously at 2% per year for the full time horizon of the model. This value is approximately in line with observed trends since 1900⁴¹, though growth rates of primary, final and useful energy have differed

slightly, due to technological progress and conversion efficiency increases, especially associated with electrification. Accordingly, growth rates have also varied somewhat over different technological eras in recent history. From around 1900 to 1940, energy system growth rates were around 1%; from 1940 to 1980 they were around 3%; then from 1980 to 2020 they were around 2%.

Consistent with these variations in growth rates historically, there is large uncertainty in the future development and growth of the energy system. In this work we do not attempt to model the many factors that are responsible for determining energy demand and supply, but instead construct exogenous scenarios that represent the net effects of all underlying supply and demand drivers on total system size. It is entirely plausible that a very different growth path from our 2% central assumption will be followed, and therefore we test the sensitivity of our main results to two alternative annual end-use useful energy growth rate assumptions: 1% and 3%. These lead to very different size energy systems over the next fifty years (2 to 3 times different in size). By considering all three growth rate cases we cover roughly the full range of system growth paths modelled in the literature (we estimate the compounded annual growth rate (CAGR) for useful energy between 2020-2070 for the five IPCC Shared-Socioeconomic Pathways to range between 0.82% (SSP4) and 1.95% (SSP5)^{207,208,209,210,211,212,213,214}).

To construct the 1% and 3% growth scenarios, we simply take the existing scenarios that are based on the 2% growth assumption, and scale the quantity of each technology in each scenario by the appropriate factor required to go from 2% annual growth to the new growth factor in year t , i.e. $(1.01/1.02)^t$ and $(1.03/1.02)^t$. The resulting scenarios are shown, *in terms of primary energy*, in Figure S51. We show primary energy here instead of useful energy because this reflects best the method used to perform the transformation. Since the conversion factors from primary to final to useful energy are all linear in our model, as are the relationships between VRE sources and storage quantities, these new scenarios correspond to annual growth of end-use useful energy services of 1% and 3%, as desired. In addition to the primary energy sources shown, all storage technologies scale in exactly the same way, so that all storage capabilities remain identical (i.e. in terms of duration and size relative to VRE sources and sectors). Bear in mind that much of the primary solar and wind electricity in the transition scenarios is converted to P2X fuels for delivery to end-users, so final and useful energy will be lower, and will include the P2X energy carrier (as shown in e.g. Figure S8).

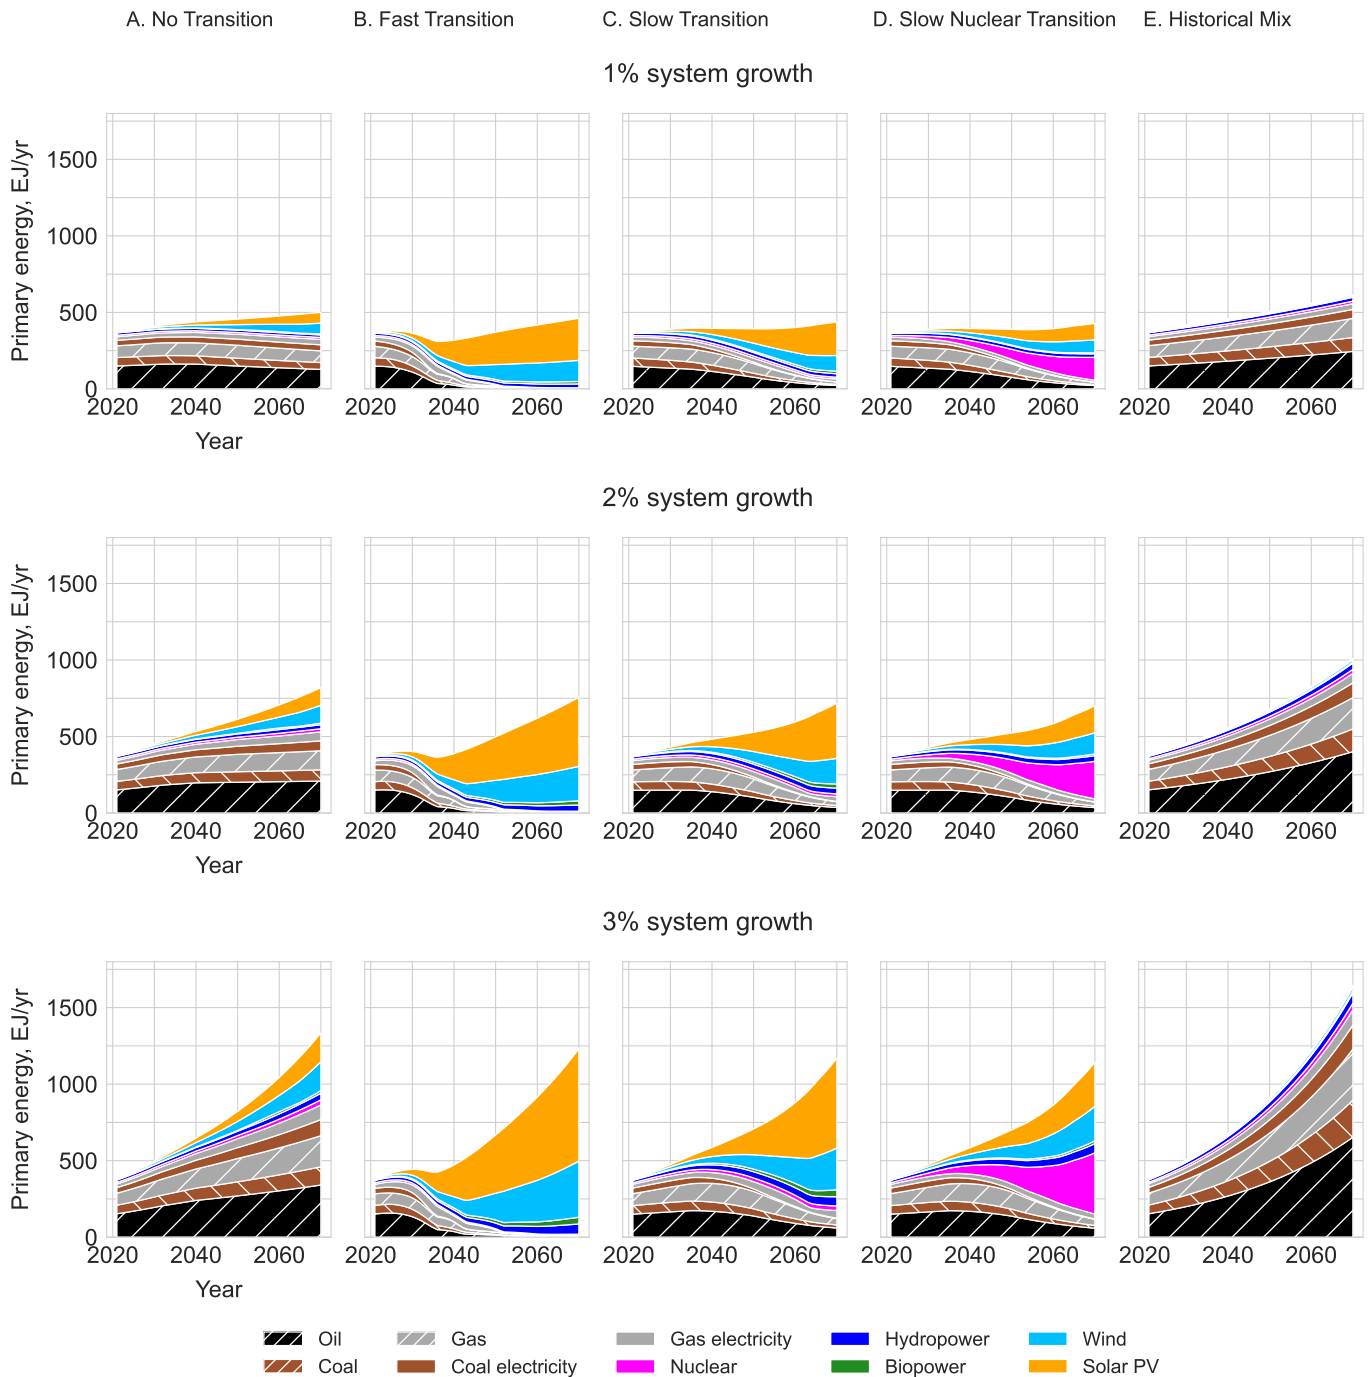


Figure S51: Primary energy in the five scenarios (from left to right), in cases where exogenous growth in useful energy services in end-use sectors is 1% (upper row), 2% (middle row) and 3% (lower row) annually.

An important feature of the way these alternative growth path scenarios are constructed is that technology growth rates are not held fixed, but are also scaled up or down as a result of the transformation. However, the effect is very small – for example, while the initial growth rate of solar PV in the Fast Transition scenario with our central 2% growth assumption is 28% (see Table S10), the corresponding initial growth rates in the 1% and 3% growth cases are around 26% and 29%. These variations are small enough that they will not significantly impact the results.*

The key unknowns that may be responsible for such different future growth paths are believed to be:

- *End-use efficiency improvements.* Well known examples are the dramatic increase in efficiency in the provision of lighting services¹⁴⁹, vehicle fuel efficiency standards, and efficiency increases in computing services²¹⁵. Such changes have the effect of decreasing useful energy demand, and are very common. They can be either primarily supply side driven, via research and innovation, or primarily demand side driven, via consumer preferences or regulations. Efficiency improvements can come about due to

*It is interesting to note that slightly higher or lower technology growth rates imply that progress along experience curves will be slightly faster or slower in the alternative system growth scenarios, though the effects appear to be negligible in this case.

technological progress either with or without energy carrier switching. Our model includes the major opportunities for efficiency improvements via carrier switching, but does not include other end-use efficiency improvements.

- *Technology costs and the rebound effect.* Contrary to the exogenous scenario construction process used in our model, the real world energy system develops as a result of the interaction between energy prices and energy demand (plus many other factors). If energy prices decrease due to falling renewables and storage costs then demand for energy services may increase. Although we do not model this process dynamically (as most other models do), by considering a wide range of system growth paths, we are able to provide appropriate scenario comparisons that represent sufficiently well the main likely eventualities. This supports the robustness of our results to a range of uncertain outcomes.
- *Increasing demand for energy services as people get richer.*
- *Increasing demand for energy services to adapt to climate change, e.g. cooling, heating, building and infrastructure resilience.*

It is important to understand that in our modelling framework the appropriate way to compare scenarios is across scenarios that deliver the same quantity of useful energy services to the economy. Even if rapidly falling electricity costs cause an increase in electricity consumption, this energy is used for productive purposes and contributes to welfare increases, so for an apples-to-apples comparison of scenario costs, we must only compare it with other scenarios that provide equal levels of energy services.

Figure S52 shows the median annual expenditures on each technology in each scenario, in each of the different exogenous system growth rate cases. Figure S53 then shows the relative net present cost of each scenario as a function of the discount rate for each case (*cf.* Figure S47). These show that our main results are robust to variations in our exogenous system growth assumption, and in fact hold for significantly different sized energy systems.

The interpretation of this robustness is that the sheer quantity of solar, wind, batteries and electrolyzers made in the first few decades of the Fast Transition scenario, in any system growth case, is large enough to allow costs to progress (stochastically) far enough down their experience curves to make the entire transition scenario worthwhile in terms of NPC, regardless of what happens later in the scenario. As with earlier results, it is important to recognise that while this may appear “obvious” to some in hindsight, these results could not have been known *a priori*, and are not simply the result of the choice of cost models. They arise by performing stochastic simulations using parameter values obtained by calibrating different empirically validated cost forecasting models to technology-specific data. Thus it is the specific combination of technology parameters emerging directly from data, in particular the experience exponent and the exogenous noise parameter, that lead to the results.




Figure S52: Median annual expenditure per technology in each scenario (left to right), for exogenous system growth rate assumptions of 1% (upper row), 2% (middle row), and 3% (lower row).




Figure S53: Expected relative NPC vs discount rate in different system growth scenarios, with time horizon 2070.

7.4 Sensitivity to cost assumptions

To explore how results depend on technology cost and uncertainty parameters, we present a few side cases reflecting alternative parameter choices: 1) Less pessimistic solar costs; 2) Lower oil and gas prices; 3) Lower nuclear costs; and 4) Less uncertain electrolyzer costs. In each case only the stated parameters are changed, all others remain identical to the main case. In the following analyses we took $M = 10000$, to reduce computation time.

7.4.1 Less pessimistic solar costs

The main case was initialised with a current cost of solar electricity of $c_0^{10} = 57$ \$/MWh. However, LCOE values from different organisations vary due to methodology and data sources. BNEF’s 2020 global values for tracking and non-tracking PV are 40 \$/MWh and 49 \$/MWh, respectively. In this side case we use $c_0^{10} = 49$ \$/MWh.

Solar costs are moving so fast it is important to use the most up to date values possible. Since the experience curve parameters remain unchanged (as these are formed from the entire history of costs, which does not change, apart from one point), changing the initial cost has the effect of just shifting the experience curve down. This will have big implications for the cost of transition, as the median and means are also shifted down, and for solar this has a large effect. The cost forecasts for solar in this side case are shown in Figure S54, and the NPC results are shown in Figure S55. At 4% discount rate the Fast Transition scenario consistently results in an expected NPC saving of around 7-10 trillion \$.




Figure S54: Solar PV electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row), when initialised with less pessimistic current cost $c_0^{10} = 49$ \$/MWh (instead of 57 \$/MWh in main case). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S35.




Figure S55: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the five scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with less pessimistic solar current costs

7.4.2 Lower oil and gas prices

Recall, the AR(1) models for oil and gas were calibrated on price data beginning in 1973 and 1985 respectively. If instead we use the entire historical data sets, back to 1861 and 1922 respectively, the estimated parameters are for oil: $\tilde{\phi} = 0.9058$, $\tilde{\sigma}_\epsilon = 0.2690$ and $\tilde{\kappa} = 0.1522$; and for gas: $\tilde{\phi} = 0.9778$, $\tilde{\sigma}_\epsilon = 0.1880$ and $\tilde{\kappa} = 0.0167$. These parameters represent prices falling back to levels seen in the early twentieth century, when demand was around a hundred times lower than current levels. The forecasts are shown in Figures S56 and S57. In this case the long term median price forecast for oil is around 30 \$/barrel and for gas is around 1.8 \$/MMBtu, i.e. global average prices are forecast to be below these levels for half the time. The NPC calculations are conditional upon such sustained very low prices. We perform the NPC calculations and present the results in Figure S58. The results are mixed, and for some combinations of discount rate and time horizon the expected NPC of No Transition is cheapest. However it is remarkable, given how low these modelled fossil fuel prices are, that the median NPC of Fast Transition remains favourable in all cases, and even the expected NPC is lowest in many.

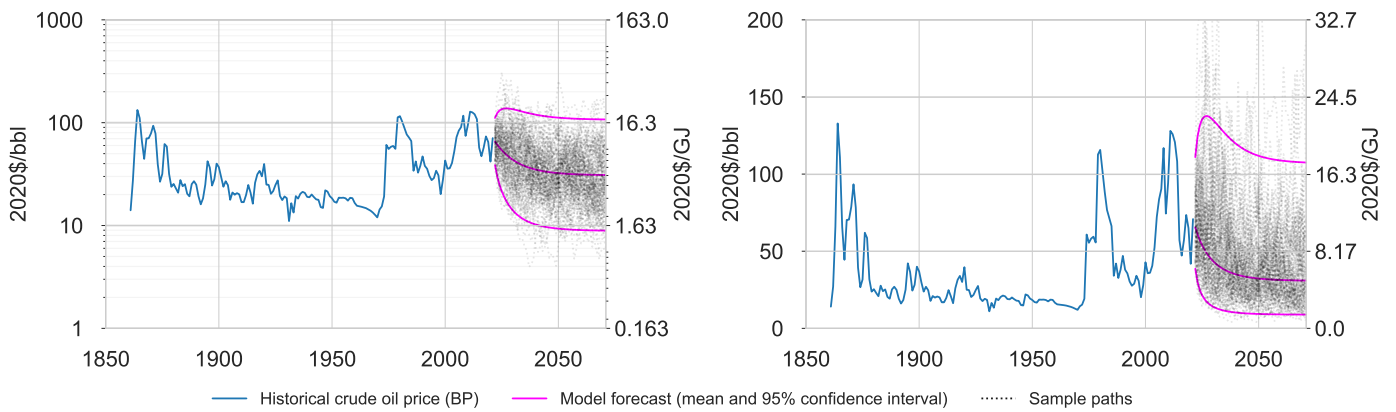


Figure S56: Low oil price forecast resulting from calibration on data 1861-2018, analogous to Fig. S18.




Figure S57: Low gas price forecast resulting from calibration on data 1922-2018, analogous to Fig. S20.




Figure S58: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the five scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with low oil and gas prices.

7.4.3 Lower nuclear costs

While the main case uses $c_0^6 = 90$ \$/MWh, which is already at the very low end of cost estimates, this side case imagines that cost levels observed in the early phase of the French nuclear rollout are witnessed on a global scale: we set $c_0^6 = 70$ \$/MWh. We also keep the very low forecast uncertainty, which, recall, was a somewhat arbitrary assumption. The cost forecasts for each scenario are shown in Figure S59 and the NPC results are shown in Figure S60. Total spend on nuclear is now much lower, making the Slow Nuclear Transition more competitive generally, but it still does not beat Fast Transition.

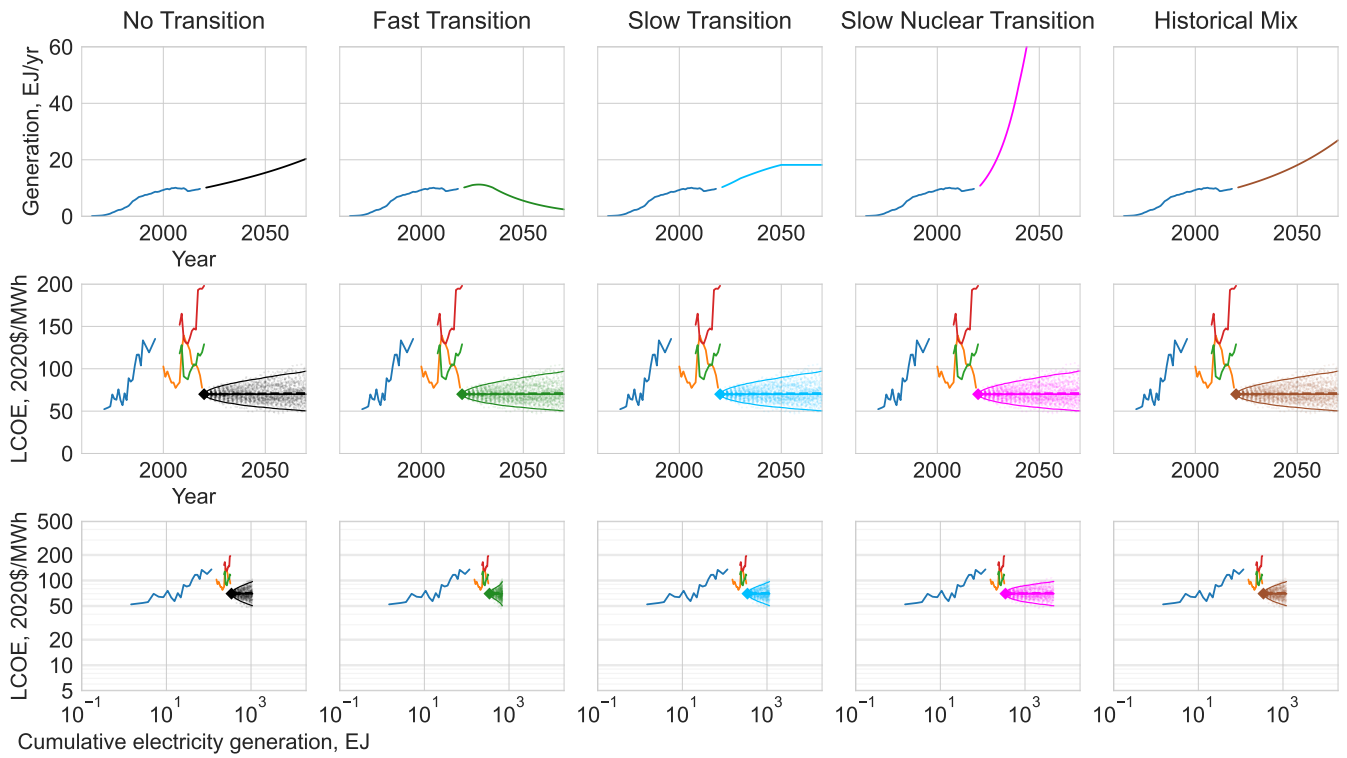


Figure S59: Nuclear electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row), when initialised with low current cost $c_0^6 = 70$ \$/MWh (instead of 90 \$/MWh in main case). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S25.




Figure S60: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the five scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with low nuclear costs.

7.4.4 Less uncertain electrolyzer costs

Based on calibrating the FDWL model on historical data, the main case uses parameters $\hat{\omega} = 0.129$, $\hat{\sigma}_{\omega} = 0.067$ and $\hat{\sigma}_{\eta} = 0.201$. These parameters give forecasts with strong median progress but large uncertainty. This is a direct result of the small sample size of the historical data. This side case considers the situation where the progress trend is the same but the uncertainty is much lower, as may occur as more data is collected and collection methods are standardised. We halve the noise standard deviation to $\hat{\sigma}_{\eta} = 0.100$. We keep the current initial cost of 1313 \$/kW. The forecasts for each scenario when these parameters are used are shown in Figure S61.

Note that in all the scenarios that deploy electrolyzers, cost forecast medians are around 300-800 \$/kW. As noted earlier, current estimates for AEL systems are already as low as 200 \$/kW, and industry expects PEM systems to reach 300-500 \$/kW within a decade, so the cost forecasts here seem quite sensible. The NPC results are in Figure S62. Now, in contrast to the main case, the expected NPCs of the Fast and Slow Transition scenarios decrease monotonically as the time horizon increases. This shows that the high uncertainty in future electrolyzer costs in the main case has a substantial effect on the results – it is precisely this uncertainty that causes the mean and median to deviate so strongly at around 2080.




Figure S61: Electrolyzer installed capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row), when initialised with lower uncertainty: $\hat{\sigma}_{\eta} = 0.100$ (instead of $\hat{\sigma}_{\eta} = 0.201$ in the main case). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figures S41 and S42.




Figure S62: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the five scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with less uncertain electrolyzer costs.

7.5 Moore's law results

An alternative theory of technological progress postulates that time is a better predictor of future costs than experience. This is known as generalized Moore's law, and forecasts made using this method were analysed in Farmer and Lafond¹³⁶. The analysis so far has used the first-difference Wright's law model given in Eq. 42; the analogous first-difference Moore's law (FDML) equation is

$$\log c_t - \log c_{t-1} = -\mu + n_t \quad \text{with IID } n_t \sim \mathcal{N}(0, \sigma_n^2), \quad (54)$$

where c_t is the cost of a given technology at time t , μ is the technology's progress rate (or Moore exponent) and n_t are periodic IID noise shocks, with technology-specific variance σ_n . As in the Wright's law case, forecasts are improved slightly by using a version of the model with autocorrelated noise shocks,

$$\log c_t - \log c_{t-1} = -\mu + v_t + \theta v_{t-1} \quad \text{with IID } v_t \sim \mathcal{N}(0, \sigma_v^2), \quad (55)$$

so we use this version. The noise shocks are related by $n_t = v_t + \theta v_{t-1}$ and therefore

$$\sigma_v^2 = \sigma_n^2 / (1 + \theta^2). \quad (56)$$

Farmer and Lafond¹³⁶ analysed the autocorrelation parameter θ for around 50 technologies and found the average was $\theta^* = 0.63$. We use this constant value of θ for all technologies in our cost forecasts here.

The parameter estimation procedure assumes $\mu \sim \mathcal{N}(\hat{\mu}, \hat{\sigma}_\mu^2)$, so when we calibrate the basic model (Eq. 54) on historical data we obtain estimates for the three values $\hat{\mu}$, $\hat{\sigma}_\mu$ and $\hat{\sigma}_n$. Eq. 56 is then used to compute $\hat{\sigma}_v$.

To generate a single cost sample path of length T , we first pick a value of μ from the distribution $\mathcal{N}(\hat{\mu}, \hat{\sigma}_\mu^2)$, then pick T successive noise shocks from the distribution $\mathcal{N}(0, \hat{\sigma}_v^2)$ and use Eq. 55 to compute c_t at each step. Repeating this sampling process a large number of times recovers the theoretical forecast distribution calculated in Farmer and Lafond¹³⁶. This process accounts for both our uncertainty in the Moore exponent (due to limited historical observations), and intrinsic future uncertainty (represented by the exogenous periodic shocks). In contrast to the Wright's law cost forecasts generated previously, Moore's law forecasts are not scenario-dependent (they do not depend on z_t), so each sample cost forecast is valid for all scenarios. We now mirror the calibration process implemented before for FDWL, in S.I. 6.7 - S.I. 6.13, but for FDML instead. Fossil fuel based technologies continue to be modeled with the AR(1) process as before.

For nuclear power we fit the FDML model to the chained series formed from the Koomey and Hultman data and the Lazard low data, which gives parameter estimates $\hat{\mu} = -0.032$, $\hat{\sigma}_\mu = 0.026$, $\hat{\sigma}_n = 0.142$. The negative Moore exponent reflects the consistent negative learning observed historically. As before though, we suppose that costs have stabilised, and use parameter values $\hat{\mu} = 0.00$, $\hat{\sigma}_\mu = 0.01$, $\hat{\sigma}_n = 0.01$. As with Wright's law, the short and inconsistent cost series available for hydropower and biopower make parameter estimation difficult. For hydropower we assume costs are stable and use parameter values $\hat{\mu} = 0.00$, $\hat{\sigma}_\mu = 0.01$, $\hat{\sigma}_n = 0.01$. For biopower we assume there is a very weak learning trend and use parameter values $\hat{\mu} = 0.01$, $\hat{\sigma}_\mu = 0.01$, $\hat{\sigma}_n = 0.01$. The FDML forecasts for nuclear, hydropower and biopower using these parameters are shown in Figure S63. This model is not well suited to forecasting beyond about 30 years for these technologies, as the forecast variance becomes very large. However, since the difference in production of these three technologies over the five scenarios is not great, the effect on the total cost calculation is dwarfed by the effects of wind and solar, as we now see.

The FDML forecasts for wind and solar are shown in Figure S64.




Figure S63: Moore's law cost forecast distributions for nuclear power (top), hydropower (middle) and biopower (bottom), on a non-log scale (L) and a log scale (R). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in S.I. 6.




Figure S64: Moore's law cost forecast distributions for solar (top) and wind (bottom), on a non-log scale (L) and a log scale (R). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in S.I. 6.

The cost forecasts obtained by using Moore’s law are much lower than those using Wright’s law. This is because with Wright’s law the cost reductions slow down as it becomes increasingly difficult to double cumulative production, but with Moore’s law the cost reductions simply continue like clockwork.

For Li-ion batteries we use the long term trend labelled “Li-ion (all applications)” on Figure S37 for calibration, yielding parameter estimates: $\hat{\mu} = 0.146$, $\hat{\sigma}_{\mu} = 0.019$, $\hat{\sigma}_n = 0.095$.

The data for Va-redox flow batteries is very sparse. The four years of Lazard low and high cost data produce Moore parameters $\{\hat{\mu} = 0.099, \hat{\sigma}_{\mu} = 0.215, \hat{\sigma}_n = 0.372\}$ and $\{\hat{\mu} = 0.233, \hat{\sigma}_{\mu} = 0.191, \hat{\sigma}_n = 0.331\}$ respectively. The Schmidt et al. data is non-sequential, so we calculate the deterministic experience exponent implied by the end points (i.e. we solve $c_t = c_0 e^{-\mu t}$), which gives a Moore exponent of 0.076. This is similar to that for wind, so even though it is lower than the Lazard data Moore exponents, we simply use this value and assume that the other two parameters are also similar to those for wind, and use these. We take $\hat{\mu} = 0.076$, $\hat{\sigma}_{\mu} = 0.015$, $\hat{\sigma}_n = 0.086$.

For electrolyzers, calibrating the model using the same PEM cost data as before yields parameter estimates: $\hat{\mu} = 0.099$, $\hat{\sigma}_{\mu} = 0.050$, $\hat{\sigma}_n = 0.200$.

The FDML forecasts for batteries and electrolyzers are shown in Figure S65.




Figure S65: Moore’s law cost forecast distributions for Li-ion batteries (top), Va-redox flow batteries (middle) and PEM electrolyzers (bottom), on a non-log scale (L) and a log scale (R). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in S.I. 6.

FDML parameters are summarised in Table S26.

Technology	Moore exponent, $\hat{\mu}$	Standard error $\hat{\sigma}_{\mu}$	Noise $\hat{\sigma}_n$
Nuclear	0.000	0.010	0.010
Hydropower	0.000	0.010	0.010
Biopower	0.010	0.010	0.010
Wind	0.083	0.015	0.086
Solar PV	0.124	0.017	0.109
Li-ion batteries	0.146	0.019	0.095
Va-redox flow batteries	0.076	0.015	0.086
PEM electrolyzers	0.099	0.050	0.200

Table S26: Summary of Moore’s law parameters.

We implemented the model using Moore’s law cost forecasts instead of Wright’s law cost forecasts, with the FDML parameters shown above, and $M = 20000$ complete ensembles. Figure S66 shows the evolution of the annual system cost distributions for the five scenarios, and Figure S67 shows the median annual spend on each technology in each scenario.




Figure S66: Annual system cost distributions, $W_t^{scenario}$, for each scenario when using the Moore’s law forecasting model. Medians are shown as solid lines, expectations as dashed lines, and the shaded regions represent the 95% confidence intervals.




Figure S67: Median annual spend on each technology in each scenario when using the Moore’s law forecasting model.

Figure S68 shows the expectations and medians of the NPC of each scenario relative to No Transition, for varying discount rates, at varying time horizons.




Figure S68: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the five scenarios, when performing the NPC calculation using different time horizons between 2040 and 2070, and the FDML.

The results show that using a Moore’s law model for technological progress makes the Fast Transition even more advantageous than when using the Wright’s law model. At all relevant discount rates and time horizons the Fast Transition scenario is expected to generate large savings compared to all other scenarios.

8 Further Discussions

8.1 The heterogeneity and persistence of technological change

Predicting the future cost of technologies is inherently difficult given the uncertainty in the evolution of technologies. However, an examination of the historical time series for a vast range of technologies shows two clear signals. Firstly, that technology costs improve at very different rates. For instance, the empirical data shown in Figure S69 provides the annual growth rates in prices for a large range of manufactured products over roughly 60 years from US producer price index data. Most technologies are in the bell-shaped component of the data on the right, around the level of 0% change in price, but there are some clear outliers that have experienced annual price declines as high as 16 percent - such as computer electronics.

The second clear trend in the empirical record for technologies is that once they become established, these trends become remarkably predictable. The precise technological solutions might not be known, but once a technology establishes itself on an improvement trajectory, it tends to stick there for a long time. Figure S70 presents the time series for a range of different technologies, showing that such tendencies appear to cluster in the specific technology groups. Chemicals for example seem to all improve at a similar slow rate, while computer hardware, shown in red, improves at a faster rate, and genomics, in pink, even faster. This is not a comprehensive and unbiased collection but nonetheless we see a pattern in cost rates of various technologies across technology types.




Figure S69: A histogram of the annual growth in prices for various U.S. manufactured products from 1958-2011 showing the majority of products have a growth rate around 0%. Data Source: National Bureau of Economic Research database (<https://www.nber.org/research/data/nber-ces-manufacturing-industry-database>)

The two most prominent empirical laws that have been applied to capture such trends are Moore's Law and Wright's Law, both of which have been shown to work well empirically on over fifty different technologies (i.e. all technologies for which sufficient data was available)^{2,216,185}. Figure S71 shows the two laws applied to four different technologies – transistors in black, photovoltaics in the red, computer hard disk drives in green, and ethanol in blue. Both Moore's and Wright's Laws appear to also work quite well for the key technologies in this current research (as shown by the data series in S.I. 6). We apply both in this study, with a focus on Wright's Law, given its implications for policy in the context of the clean energy transition. See the discussion in S.I. 8.3.1 on causality regarding these implications.




Figure S70: The exponential decline in costs through time of a range of products by category. Note that the label “Hardware” is short for Computer Hardware. Source: Lafond et al. ².

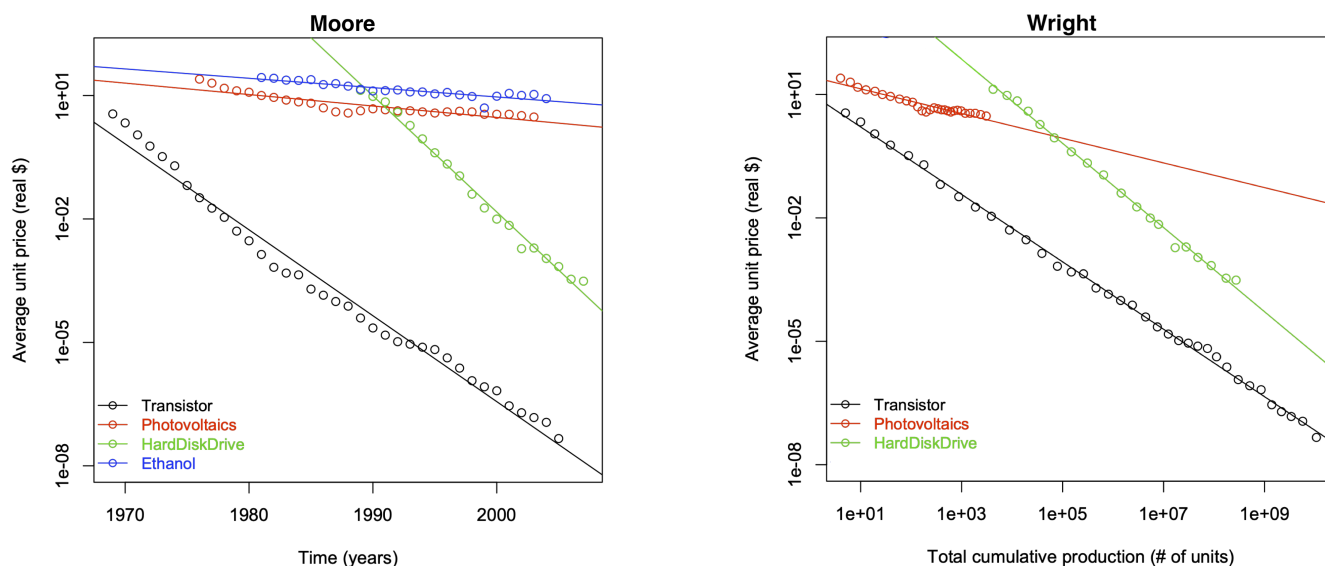


Figure S71: The application of Moore’s Law and Wright’s Law to four different technologies, demonstrating the consistency in their cost trends over time and over cumulative production. Note that while the simple trend extrapolations shown here result from the deterministic application of the two laws, the probabilistic forecasting model we use is more sophisticated – it uses the data to generate forecast *distributions* (which have been validated empirically) so that technology costs are not “guaranteed” to continue on their current trajectories, as suggested by the these plots. Source: Performance Curve Database, Santa Fe Institute (<http://pcdb.santafe.edu/>) with thanks to Bela Nagy

8.2 Bias-variance trade-off

When choosing the best model for prediction it is essential to be aware of the bias-variance trade-off. When the correct model is not known *a priori*, as is usually the case, it is necessary to use a family of models, with free parameters that are fit to the data to select the best model. There are two competing factors that cause prediction errors. Bias occurs when the family of models does not contain the true model. For example, if the true model is nonlinear, then a fit based on a linear family of models will make prediction errors. In contrast, *variance* occurs when the data are noisy and there is not enough data to fit the parameters correctly. The bias-variance trade-off refers to the fact that to make good predictions it is necessary to find the best compromise between bias and variance. This can be counterintuitive: For example, even if the true model is nonlinear, with limited data a linear model may make better predictions. Although one might naively think

that adding more factors so that the family of models includes the best model is always the best procedure, this is often not the case.

In the context of technology prediction, Nagy et al.¹⁸⁵ tested the prediction accuracy for a variety of different functional forms, and found that the best predictions were made by very simple families of models with only two parameters. In particular, Wright's law and Moore's law produced the best forecasts. Nagy et al.¹⁸⁵ also tested a combined form that included Wright's Law and Moore's Law as limiting cases, and found that this produced far worse forecasts than either of them alone. This is a good example where adding additional free parameters reduces bias but causes greater additional variance, which makes the resulting predictions worse (in this case substantially worse).

Thus, while it might be tempting to impose more complicated models for forecasting technological progress, this must be done with care. Two factor learning curves might allow better predictions, but until their ability to make good out-of-sample forecasts is properly tested, they should be used with caution. More generally, simply fitting functional forms to data is insufficient – it is essential to perform out-of-sample testing. The simple model forms that we use here have been extensively tested by making out-of-sample predictions, and the method for estimating their errors has also been extensively tested by making out-of-sample predictions.

8.3 Wright's law caveats

It is well known that experience curves are approximate, empirically observed phenomena. The underlying causal mechanisms are not well understood, which has led to criticism and a reluctance to apply experience curve concepts in practice, despite abundant empirical evidence of their existence as an emergent property of technology learning systems. This in turn has led to recent high profile prediction failures in energy modelling (for example large errors in solar PV deployment forecasts). Recent work has made progress on understanding causality in learning systems though, and has shown that experience is indeed a prominent causal factor in cost reductions in many cases²¹⁷. We discuss this here in more detail given its relevance to this present research.

8.3.1 Discussion on questions of causality

In using Wright's law there is a fundamental question about causality. Does increased cumulative production cause prices to drop? Or alternatively, does production increase because prices drop? One can make plausible arguments both ways.

The argument that decreasing costs cause production to increase is simple: Decreased costs cause increasing demand, which causes increasing production. The argument that increasing cumulative production causes costs to drop is a bit more complicated. It is based on the idea that production is proxy for effort. More production means that more units are made, which provides more opportunities for learning. Increasing production can also cause more competition and can lead to economies of scale. Economies of scale are explicitly present in chemical production, for example, where the cost of production may be dominated by the material cost of building a plant; since the volume of a vessel of a given diameter increases faster than the area of the material required to build the vessel, costs drop automatically as the vessel gets bigger. Economies of scale can enter in many other ways, e.g. increased production can lower unit costs for R&D, manufacture, marketing, sales and distribution.

In reality, of course, causality flows in both directions. There is a virtuous cycle, in which increasing production causes lower costs and lower costs cause increasing demand, which increases production. To make this clearer we present a mathematical model below. This model is based on Nordhaus¹²⁶ and is copied from an exposition given in Lafond et al.²¹⁷. It makes these issues mathematically precise and makes it clear how the underlying parameters that determine the interaction between supply and demand are reflected in an empirical learning curve. We then briefly review a study of how costs changed during World War II, which can be viewed as a natural experiment in which the scale of production was caused by wartime needs rather than cost. This study provides evidence that production causes decreasing costs (which does not rule out causality in the other direction).

Let Q_t denote production at time t and let Z_t denote cumulative production. Assume the unit cost function c_t depends on both experience Z_t and an exogenous time trend, with the functional form

$$c_t = c_0 Z_t^{-b} e^{-at}. \quad (57)$$

The parameters a and b are expected to be positive, since we generally observe technological progress. Note, however, the difference in the cause of progress corresponding to the two terms on the right. The term Z_t^{-b}

corresponds to Wright's law; Z_t corresponds to the production of a specific technology, and in general we expect b to be technology specific. The term e^{-at} , in contrast, corresponds to a generalized form of Moore's law. This can be viewed as an inexorable force that causes costs to drop with time, regardless of the level of production or any other causal factors. If a is the same for all technologies, then this corresponds to technological progress in general. It is also possible that a is technology specific.

Assume that production is equal to demand and price is equal to unit cost. Furthermore assume that demand has constant elasticity $\epsilon > 0$ and that there is an exogenous and growing demand e^{dt} . This gives*

$$Q_t = D_t = D_0 c_t^{-\epsilon} e^{dt}. \quad (58)$$

By letting Δ represent the time difference operator $\Delta x = x_t - x_{t-1}$, Eqs. 57 and 58 can be written as

$$\Delta \log c = -a - b \Delta \log Z \quad (59)$$

and

$$\Delta \log Q = -\epsilon \Delta \log c + d. \quad (60)$$

Now suppose that production grows exponentially. Then experience also grows exponentially at the same rate, that is

$$\Delta \log Q \approx \Delta \log Z.$$

Using this in Eq. 60, the solution of the system Eqs. 59-60 is

$$\Delta \log c = \frac{-a - bd}{1 - b\epsilon}, \quad (61)$$

$$\Delta \log Z = \Delta \log Q = \frac{a\epsilon + d}{1 - b\epsilon}. \quad (62)$$

Nordhaus¹²⁶'s critique is that experience curve studies typically assume that the logarithm of experience is the single explanatory factor for the logarithm of costs, i.e. they assume that

$$\Delta \log c = \beta \Delta \log Z, \quad (63)$$

and perform a regression of the logarithm of costs against the logarithm of experience in order to measure the empirical parameter β . Inserting Eqs. 61 and 62 into Eq. 63, we can write the parameter β as the ratio

$$\beta = \frac{\Delta \log c}{\Delta \log Z} = \frac{-a - bd}{a\epsilon + d}. \quad (64)$$

It is clear that in general the empirically estimated parameter β not only depends on the experience parameter $-b$, but it also depends on the exponential trend parameter a for costs, and the exponential trend parameter d for exogenous demand growth, and the demand elasticity ϵ . The estimated β can be interpreted as being purely the effect of experience only when there is no exogenous time trend, i.e. if $a = 0$ then $\beta = -b$.

Lafond et al.²¹⁷ collected an extensive data set on costs and production for the U.S. military effort during World War II, and used this to study the relative importance of increasing production vs. overall technological progress as causes of decreasing costs. Their central argument is that during World War II production decisions were made to "win the war at all costs". Roosevelt did not say, "Tanks are getting cheaper, let's build more", but rather ordered the production of as many tanks as were needed to win the war.

In this case causality is clear and it becomes possible to distinguish the effect of an exogenous trend from that of increasing production. If demand is completely exogenous then $\epsilon = 0$. If we assume for simplicity that demand is still growing exponentially, then Eqs. 61 and 62 become

$$\Delta \log c = -a - bd, \quad (65)$$

and

$$\Delta \log Z \approx \Delta \log Q = d, \quad (66)$$

so that putting Eq. 66 into 65 gives

$$\Delta \log c = -a - b \Delta \log Z, \quad (67)$$

*Nordhaus¹²⁶ also includes population growth in the exogenous demand term.

which is the same as Eq. 59. If demand is completely exogenous, we can estimate the supply equation directly[†]. A simple first-difference regression model for the logarithm of cost against the logarithm of experience with an additive constant is able to separate the effect of the exogenous trend parameter a and the “learning” parameter b .

U.S. military production during World War II is a particularly good case study for another reason. If production increases exponentially then cumulative production also increases exponentially. If in addition costs decrease exponentially, then by eliminating time from the two relations and solving for costs in terms of cumulative production, one trivially gets Wright’s law. Most of the data sets that are available show both decreasing costs and exponentially increasing production, which makes it difficult to distinguish an exogenous time trend from Wright’s law.[‡] In WWII, in contrast, the U.S. went from the 18th largest military in the world to producing 2/3 of the military equipment used by both sides combined, and then dramatically decreased production as the war ended. This means that production is far from an increasing exponential, which makes it much easier to identify the two terms. The result of this study indicated a roughly 50-50 split between production specific effects and an overall trend. From the point of view of our use of Wright’s law here, this is good news, as it indicates a substantial causal effect of cumulative production on costs, which is all we need.

In fact, since our main goal here is to determine whether a fast renewable transition is likely to be cheaper than alternative scenarios, the issue of causality is not important. Our Fast Transition scenario simply extrapolates existing production growth trends. This means that all that is needed is for the existing feedback relationship between supply and demand to continue. While these trends continue, it doesn’t matter whether one uses Wright’s law or Moore’s law – the answer will be the same. They only become different once the S-curve reaches saturation due to renewables becoming dominant. When this happens the current trend, in which renewables growth rates are much higher than overall energy growth rates, is broken. After this happens Wright’s law predicts that costs should drop more slowly, at a rate that is governed by overall energy production. Moore’s law, in contrast, predicts that costs will continue to drop at the same rate as they are dropping now. This means that Moore’s law results in cost forecasts for renewables that are more optimistic than those we make here. If this is true then our arguments here only become stronger. This is demonstrated explicitly in S.I. 7.5.

8.3.2 Further caveats

While it is critical to ensure that experience curve patterns of technology development are represented in any model, some additional limitations of the framework must also be understood. Well known caveats and subtleties include:

- *Technology function, boundaries and aggregation.* Technological improvements are the result of many different factors, including improvements through experimentation, manufacturing, modularity, economies of scale, deployment, operational characteristics, market dominance, user feedback, and increased confidence in financial returns. As a result of all these effects, technologies are constantly changing and evolving. The question therefore arises whether each new version of a technology should be treated as having its own experience curve. Is each new version of a wind turbine a new technology, or is it a continuation of an existing design? Are the varying designs from different manufacturers of wind turbines different technologies? Modelling choices must be made regarding how to aggregate such technologies in to a single overarching group forming an independent experience curve. Properties to be considered include technology type (concurrent models), vintage (chronologically sequential models), function, and geographical location (area over which the learning system extends). These choices impact experience curves, learning rates and variability observed. In this paper we use high levels of technology aggregation and global geographical scope in order to capture major trends in the global energy system clearly and effectively. We have made careful choices of which technologies and aggregation levels to use based on their function, substitutability, historical data record and relevance within the energy system.

[†]Pozzi and Schivardi²¹⁸ used survey estimates of demand elasticity coupled with data on firm-level prices to separate the effects of demand and productivity on firm growth.

[‡]While illustrative of the reverse causality issue, this model is not stochastic and all variables grow exponentially, so it cannot do justice to the other important issue, which is multicollinearity. If production grows exponentially with fluctuations, since computing experience from production involves integrating and therefore smoothing, Z will increase exponentially but with much lower fluctuations, so the term $\Delta \log Z_t$ will be approximately equal to the constant d , for all t . To give a concrete example, if production follows a geometric random walk with drift r and noise σ , then the change in the log of cumulative production has a variance approximately equal to $\sigma^2 \tanh(r/2)^2$, which since r is small is much smaller than σ^2 .

This aggregation allows us to observe relatively long experience curves historically, over timescales relevant for the energy transition. For example, we choose “wind” as an aggregate rather than separating out “onshore” and “offshore” wind, as both share the same designers, manufacturers, and technologies and learn from each other’s experience. We also use “solar PV” as an aggregate rather than separate out the many versions of PV technologies available, such as silicon solar PV, organic solar, and perovskites. As the currently dominant design for solar PV, crystalline silicon, provides the majority of the empirical data for our probabilistic forecasts, it could be argued that our aggregate solar PV forecasts are conservative in the long run as this technology is likely to be surpassed by superior PV technologies over the timescales that we model.

- *Technology surprises.* The future is uncertain and technologies will always be invented that render past technologies obsolete, we just don’t know how or when these will occur. This is a problem faced by every technology modelling method, including all existing energy models, not just experience curve based methods. To mitigate the risk of a future technology surprise rendering ongoing R&D, investment and deployment expenditures less valuable than expected, we have chosen experience curve technology boundaries carefully, considered potential technology substitutions (based on physics and engineering constraints), and obtained the best possible estimates of future cost reduction probabilities. Historical data shows that technologies tend to progress steadily and systematically, and that many events that are considered surprising by some (such as recent cost reductions in solar PV and batteries) are in fact just part of stable long-term trends. Another important consideration regarding technology surprises is that to have a major effect on mitigating climate change a technology must reach global scale within a few decades at most. This would be a huge challenge for any new technology and so, all things being equal, the risk of high impact technology surprises occurring in our problem domain is slightly lower than in a system without such a constraint.
- *Technology lock-in.* Closely related to the risk of technology surprises is the risk that by providing high support to one technology, a better technology is deprived of the opportunity to develop sufficiently, and the inferior technology becomes locked-in. Again, this issue is addressed (as best it can be) by our judicious choices of technology aggregation. For example, we are not overly prescriptive about which type of batteries or solar cells are used, we assume that many competing options may exist simultaneously as innovation and experimentation within each class of technology progresses.
- *Technology failure.* Technology costs will not fall exactly as anticipated by a deterministic experience curve model. Thus backing one technology entirely can result in wasted effort, and it is often better to back some combination of substitute technologies. This problem is addressed by our use of distributional experience curve cost forecasts. When using this probabilistic method, portfolio analysis can be applied to the decision space to determine optimal combinations of technology production for any given level of risk aversion¹²⁷. In practise this method can only be applied to small regions of decision space because non-linearity of the experience curve produces a non-convex objective function that prevents the use of linear optimizers to find a global solution. We use a variant of this approach in this paper – the decision space is too large to explore fully, so instead of attempting to optimise technology portfolios over any region of decision space, we construct five very distinct technology portfolios (i.e. scenarios) and order the resulting cost distributions.
- *Experience curve decomposition.* Many mechanisms contribute to a given technology’s innovation system, e.g. invention, research, development, demonstration, niche markets, investment, production, economies of scale, materials improvements, efficiency improvements, spillovers, deployment, learning-by-using. It is not known in detail if, how or to what extent different aspects of the system produce progress along the experience curve. In particular it is of great interest whether R&D or deployment results in greater cost reductions. This is an important question, and many studies have attempted to decompose experience curves in to multiple factors, but we do not address this issue here. We instead choose high levels of technology aggregation and assume that all effects proceed concurrently. We assume that the existing ratio of investment into R&D and deployment remains unchanged for roughly the next decade.

8.4 Comparison to existing energy system modelling approaches

There are a plethora of energy system models that have been designed to analyse energy-system planning, environmental impacts and energy-economy interactions. Depending on their intended use, these models vary

in their structure (e.g. top-down vs bottom-up), modelling techniques, sectoral, and geographical coverage and technological detail^{219,220}. Many models (e.g. MESSAGE²²¹, MARKAL/TIMES^{222,223}, OseMOSYS²²⁴) assume perfect foresight and are designed to optimise a given energy system from a social planning perspective based on a number of constraints (e.g. to enforce system performance criteria or environmental considerations)²²⁵. However, there are numerous other models employing different modelling approaches, such as POLES²²⁶, which is based on a partial-equilibrium, dynamic-recursive simulation and LEAP²²⁷, which employs an accounting-based framework.

The model developed in this paper differs from such models in a number of key ways. First, *its simplicity*: in response to what has been described as a “crowded landscape of model-based analyses that can overwhelm decision makers with their complexity”²²⁵, we have deliberately aimed to develop a novel energy system modelling approach that is as parsimonious, transparent and empirically grounded as possible. We consequently only represent energy system components that are necessary to model the key drivers that will determine how the cost of the global energy system is likely to change under different scenarios. We also limit our choice of technologies and scenarios to those whose costs can be appropriately forecast with robust historical data.

Second, *our choice of scenarios to compare*: most traditional energy system models construct a reference or Business-As-Usual (BAU) (or No Policy) scenario as a hypothetical future in which no policy interventions are assumed to occur, and compare this to alternative scenarios, which usually include some policy interventions (such as the introduction of a carbon tax or a subsidy scheme). This can result in very different scenarios being grouped together in the same analysis (e.g. scenarios with very different energy service consumption levels), and perhaps inappropriate comparisons being drawn. Merely constructing a reference scenario in this way requires many strong assumptions about population growth, economic growth, climate impacts, technology costs, technological progress and innovation, and the role of fossil fuels in society²²⁸. The choices made tend to result in framing the decarbonisation of the energy system as a cost penalty compared to continuing along some proposed “normal” fossil fuel trajectory²²⁹. This also implicitly makes strong assumptions about what the market alone is likely to deliver relative to the “social planner” perspective – neither of which are particularly realistic future scenarios²³⁰. As our model is designed to investigate future energy system transition scenarios by extrapolating certain subsets of existing technology deployment trends, we instead construct a reference scenario simply by assuming that the current technological structure continues into the future. We then compare this constant growth scenario to alternative scenarios that involve changes in the underlying technological structure of the global energy system, but which deliver equivalent levels of energy services throughout the economy. We believe that this is the most suitable approach for addressing the question of “how much it will cost” to transition from the energy system as it is currently structured to one based on clean energy technologies.

Third, *our focus on technology costs*: given the importance that learning plays in determining technology costs, many energy system models have been adapted to incorporate endogenous technological learning (ETL) whereby technology-specific investment costs decline as a non-linear function of the technology’s cumulative capacity^{231,232}. However, as most energy system models rely on linear programming, they generally need to approximate learning curves using linear segments and binary variables (an approach known as mixed integer programming), which can be very computationally costly²²⁵. While more recent efforts have reduced computing time by employing data clustering approaches²³³, resulting scenarios tend to be highly influenced by the assumed capacity build-rates or floor-costs, which appear to have limited empirical basis. These constraints are often exogenously input into models to avoid unrealistic “tipping points” where all investment is directed to a given technology once it becomes marginally cheaper²²⁹. Such practices have led to open questions about whether ETL is worth the trouble and additional computational burden, given that results are ultimately largely conditioned by these exogenously imposed limits²³⁴.

The technology forecasting approach we describe in this paper provides an alternative methodology that has the advantage of being simple and transparent. Rather than relying on least-cost optimisation approaches (which have been shown not to reflect real-world energy transitions²³⁰), we exogenously construct scenarios using a few simple rules constraining the supply, growth and substitutability of technologies, then use statistically validated forecasting models to calculate probability distributions of technology and energy system costs, as a function of these scenarios.

It is beyond the scope of this paper to attempt to solve for the least-cost scenario. In fact, optimizing over the space of scenarios is an intrinsically difficult problem due to positive feedback between cost and production in the mathematical formulation of the experience curve. This feedback results in a non-convex objective function, so global optimization tools cannot be used. However, in a probabilistic experience curve setting, heuristics may be available, as there are both tipping points and large regions of parameter space that

are near-optimal¹²⁷. We therefore restrict our attention to exploring a few representative areas of the space of scenarios.

8.5 The use of floor costs in endogenous technological learning models

Floor costs are commonly employed in endogenous technological learning (ETL) models (Table S21) including for clean technologies such as solar PV, wind and EVs, in the Integrated Assessment Models IMACLIM, MESSAGE, MERGE-ETL, ETSAP-TIAM, POLES, GCAM, REMIND, WITCH and TIMES. There is however a general absence of explanation in the literature for how such floor costs are determined. In references provided in the associated manuscripts we were unable to find any hard empirical evidence to support the values chosen for floor costs in the associated model runs. It is not always stated but it is likely that in most cases such floor costs “have been estimated, based on own expert judgement”¹⁶⁶.

As can be seen in Table S21 the values used for PV floor costs have declined for the more recent model runs of the same IAM, presumably after the previous floor costs were shown to be wrong, but again with little explanation as to how the new values were determined. It should be noted that these floor costs are also commonly applied such that the technology investment costs asymptotically approach the exogenously prescribed floor costs^{166,179} meaning that the impact of the floor cost on deployment might be felt well before the floor cost is reached. A few studies have applied sensitivity analyses to these floor costs and shown quite clearly how important they are to the results^{182,235}.

One explanation provided for the use of floor costs is to prevent costs from tending to zero in the long run²³⁶, or in other words to reflect that “learning slows down as a technology matures”²³⁵. However, there is no obvious evidence of this to be found in the more than 100 technologies included in the Performance Curve Database¹⁸⁵. In the IAMs that couple ETL energy models with computable general equilibrium economy models limiters such as floor costs are sometimes used to avoid the generation of increasing returns to scale effects, which can cause “resolution difficulties” (which presumably means multiple optima)²³⁷.

In practice the problem for ETL-based models is that they will have a tendency to invest heavily in technologies with high learning rates unless they are prevented from doing so by additional constraints – such as floor costs²³⁴. How heavily such ETL models invest in a technology will depend on other factors including their capital to operating cost ratios and discount rates, limitations from input materials or land availability, the demand for the technology, and competing technologies. However, for a technology like solar PV, which has a historically high and persistent learning rate, few insurmountable global material or land availability barriers, and which offers a clean and cheap alternative to fossil fuels, a floor cost might be the only way to prevent solar PV from eventually dominating the global energy system in any such ETL based energy-economy model.

A possibly more defensible approach that has been used by ETL modellers is to split investment costs for a technology into two parts, one that is reducible through learning, and a second floor cost component. Using separate floor costs for the fuels for some power generation technologies has been shown to be empirically defensible. In evaluating the cost declines in coal-fired electricity generation McNerney et al.¹⁴⁰ differentiated the learning rates that could be applied to electricity generation in power plants, and the costs associated with the commodities, such as raw material or primary agricultural product – in their case, coal. Such commodities they argue are most simply modelled using a random-walk or first order autoregressive models, similar to how we have modelled oil, coal, and gas in this manuscript.

Unfortunately, the IAM literature’s use of this approach does not appear to provide any such peer-reviewed methodology for the floor cost components (e.g. Pietzcker et al.²³⁵). In some cases the assumed floor costs are documented as being based on components of the technology’s power system, such as “construction, electric facilities, road for access, etc” for onshore wind power, and “cost for power conditioner” for solar PV²³⁸. It is difficult to justify such examples. Power conditioners are quite similar to inverters, which exhibit learning behaviour²³⁹, so are also likely to experience learning effects, and as noted by McNerney et al.¹⁴⁰, although solar photovoltaics depend in part on commodities, such as metals and silicon, processing dominates the cost of mono and polycrystalline silicon, rather than commodity input costs. Solar PV therefore behaves much more like a technology than a commodity in terms of learning rates.

Finally, one example of a clean energy technology for which a floor cost might be defensible is biomass (or indeed biofuels). This is because there is a limiting factor of production (in the form of arable land), and there is an alternative use of the commodity as food (which is considered a “commons” in economic terms²⁴⁰), so that biomass has a well developed global market. Insofar as the profits from growing food remain constant through time, this introduces an effective floor cost, as market incentives will push the owners of arable land toward growing food if the cost of energy is too low. This is, however, an exceptional case, and we can think

of no other such arguments that would introduce floor costs for the technologies that we rely on in the Fast Transition scenario.

8.6 Regional differences in competing technologies

The probabilistic estimates of future costs for electricity generation technologies used in this research are based on the global average levelized cost of electricity for each technology as reported in each data source. For instance, we largely rely on global average LCOEs provided by IRENA for solar PV and onshore wind, for our solar PV and wind historic trends respectively. That said, there are of course important regional differences in costs. We provide examples in Figure S72, which depicts regional variation in photovoltaic solar energy, and Figure S73, which depicts regional variation for wind. In each of these figures we show the historical geographic variation in costs from 2010 to 2019, and then show a simple projection of these regional variations going forward. As we argue below, while regional variation is important, valuable insights can be gained when regional trends are examined through the lens of global trends.




Figure S72: Past (2010-2019) and future (2020-2070) projections of the 5th and 95th percentile global solar PV LCOE in the Fast Transition. Past 5th and 95th percentiles are provided by IRENA¹⁵⁹. Future projections of these percentiles were generated by applying the cost reductions associated with global cumulative deployment for solar PV to recent (2019) 5th and 95th percentile costs for solar PV.

The variation in levelized costs of electricity for variable renewable energy technologies by location are generally dominated by local resource conditions (solar irradiance and wind patterns), and local operating costs, including labour costs²⁴¹ and interest rates²⁴². Other factors include plant design, technology, location relative to the grid, and local market component pricing¹⁵⁹, all of which will impact on the capital and operating costs, as well as the capacity factor for any given installation.

It would be reasonable to assume that most of the global deployment to date would therefore have occurred in those locations where these factors are most favourable (low interest rates, high winds and solar irradiance, low labour costs), thereby biasing our empirical record. Definitely, within individual states, provinces, and even countries, the average will be skewed towards sites with the best solar irradiance and wind^{243,244}. However, at the global scale, the current levels of deployment of renewables do not match well with renewable potential, suggesting political will might also be a strong driver of renewables deployment²⁴⁵.

The countries with the largest installed variable renewable capacity in the IRENA database are China, the United States, India, and Germany¹³⁹. China and India have historically had comparatively low labour costs, while favourable interest rates enabled Germany's heavy investment in capital intensive solar PV during the early 2000's. None of these countries have the best solar irradiance or wind potential, but they do have some of the lowest solar PV installation costs in the world^{246,102}. One explanation for this is gains they have




Figure S73: Past (2010-2019) and future (2020-2070) projections of the 5th and 95th percentile global wind LCOE in the Fast Transition. Past 5th and 95th percentiles are provided by IRENA¹⁵⁹. Future projections of these percentiles were generated by applying the cost reductions associated with global cumulative deployment for wind to recent (2019) 5th and 95th percentile costs for wind.

made through learning-by-doing on installation which can decrease costs of system design, customer acquisition, cabling/wiring, monitoring and financing. This means that, in terms of renewable potential, many regions such as Africa, Australia, and the Middle East for solar (Figure S74), and Russia, Canada, Greenland, Argentina, and New Zealand for wind^{247,248,245} (Figure S75), have yet to capitalise on their abundant and relatively inexpensive renewable energy resources.

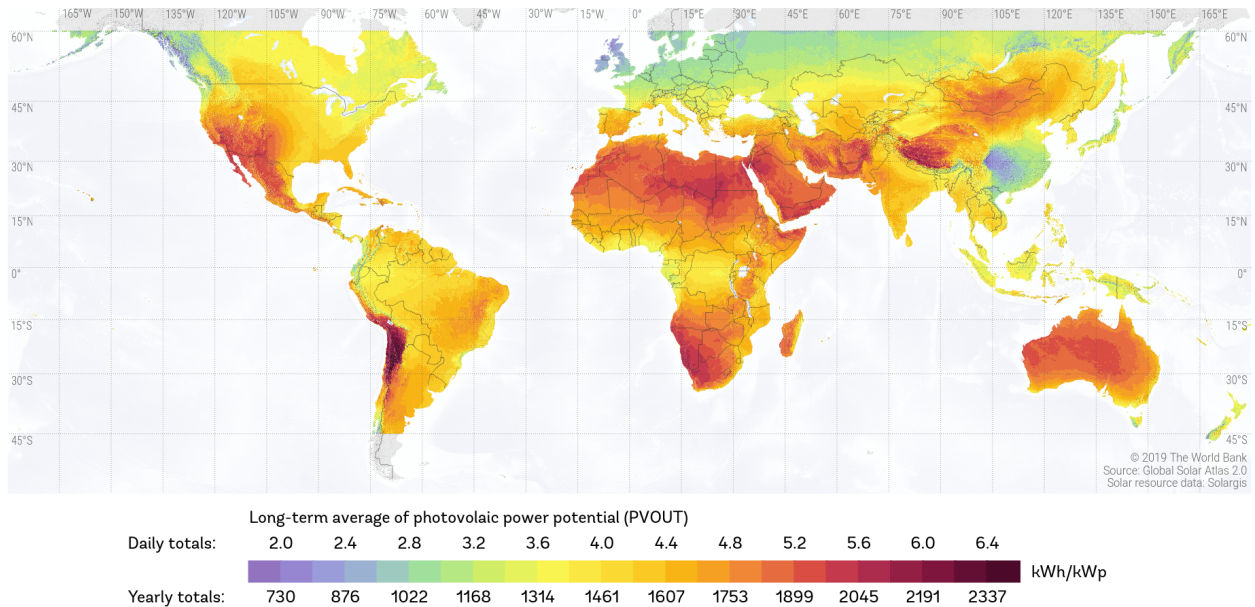


Figure S74: A map of estimated solar PV power generation potential expressed as kWh/kWp (kWp means “kilowatt potential” here). It represents the long-term average of yearly/daily potential electricity production from a 1 kW-peak grid-connected solar PV power plant. Source: solargis.com.

As costs decline even less favourable locations become cost-effective and are used with the improved technology. Meanwhile, the best sites in a region eventually carry old technologies. When these installations are retired they are replaced by newer versions which can make better use of the site’s good resources, e.g.

new, larger turbines are leading to greater capacity factors and hence lower overall LCOEs¹³⁹.

Finally, in regions where costs are currently high, such as Russia and Japan, it is reasonable to assume that without significant policy support new solar and wind deployment will continue to be slow. However, even in regions where solar and wind deployment is slower than the global average, the existence of global markets and global learning still enables costs to decline in such low deployment regions. If we therefore assume global experience influences the costs in all regions equally (i.e., global cumulative deployment also drives down costs in the 5th and 9th percentile locations), the costs for solar PV in the 95th percentile (most expensive) locations on earth are still projected to be cheaper than the least expensive global coal electricity by 2033 and the least expensive global gas electricity a decade after that (Figure S72). Most importantly, given the rate of global cost declines it only takes a few years for the regions with the 95th percentile costs to drop to the previous global average, and a decade for them to fall to match the costs associated with the countries that were in the 5th percentile. For instance, the 5th percentile LCOE in 2010 was 188 \$(2019)/MWh compared to the 95th percentile in 2019, which was 190 \$(2019)/MWh.




Figure S75: A map of global wind power density potential expressed as W/m^2 at 100m height. Source: Global Wind Atlas 3.0, a free, web-based application developed, owned and operated by the Technical University of Denmark (DTU). The Global Wind Atlas 3.0 is released in partnership with the World Bank Group, utilizing data provided by Vortex, using funding provided by the Energy Sector Management Assistance Program (ESMAP). For additional information: <https://globalwindatlas.info>

We can see from Figure S76a the extent to which the entire distribution of costs have decreased relatively evenly for solar PV from 2010 to 2019. The same can not be said for wind, with the cheaper locations seeing a greater share of the cost declines in Figure S76b. The tail appears to be lagging. If wind follows the Fast Transition cost forecast median, it will still be cheaper than the least expensive global coal by 2035, but not necessarily cheaper than the least expensive global gas before 2070 (Figure S73). Therefore, if it proves too expensive to transport clean energy to locations with good fossil fuel reserves and poor wind and solar resources, it is possible that a scenario in which some coal and gas are still used in 2070 in such regional pockets might still be cheaper than the Fast Transition. However the vast majority of the global energy system in such a scenario will consist of solar and wind generation, with associated storage.

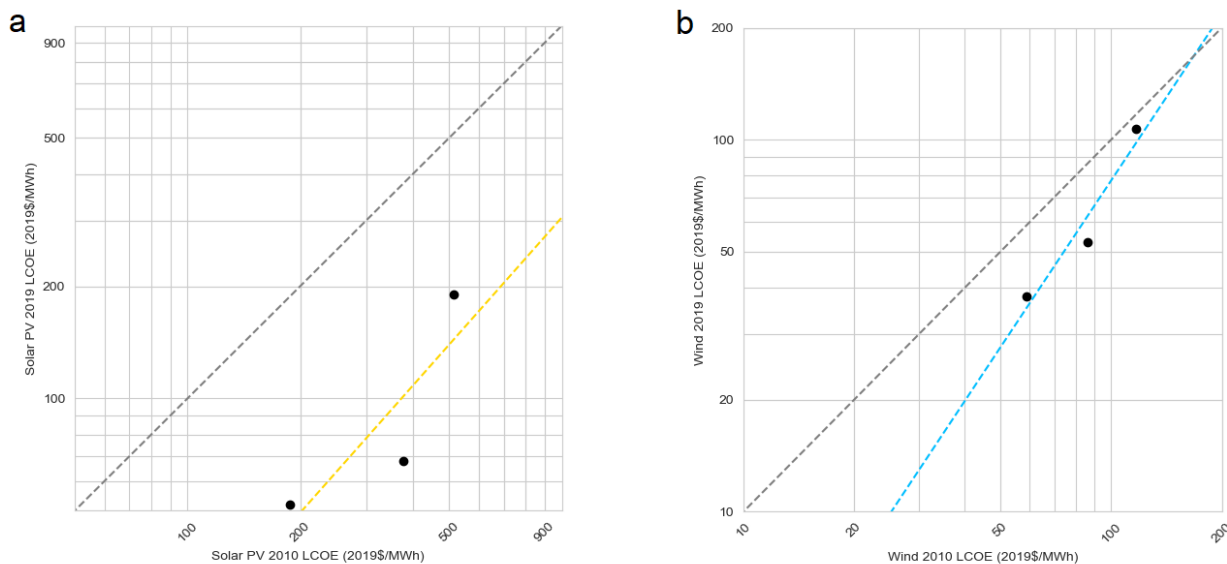


Figure S76: A scatter plot comparing the 2010 LCOE with the 2019 LCOE mean, 5th percentile and 95th percentile respectively for a) Solar PV and b) Wind (in 2019\$) including a trendline for solar PV (gold) and wind (blue) and an identity line (gray) for each.

8.7 Is the speed of the Fast Transition achievable?

The approach presented here provides a prediction of the costs associated with various scenarios of exogenous deployment of clean technologies. It is important to understand that it is the costs that are forecast, not the rates of deployment. The rates of deployment are chosen for each scenario with the Fast Transition based on a continuation of current deployment growth trends in key clean technologies, principally solar, wind, batteries and electrolyzers. We have endeavored to ensure that this scenario is consistent and achievable from an engineering perspective but this does not mean it is inevitable or even necessarily likely.

There can be little doubt that even if the Fast Transition is cheaper to engineer than the Slow Transition, significant political support will still be required to ensure continued high growth rates of deployment of renewables and storage technologies. If the drivers of existing trends do not persist into the future (e.g., generous price support for renewables, manufacturing scale economies in production, and favourable permitting and other balance of plant costs), then the cost trends we have forecast for these high deployment growth rates might not occur.

Proponents of the socio-technical transitions (STT) perspective have argued that most energy-economic models used for constructing plausible transition pathways fail to accommodate all the elements vital to understanding the pace and cost of a transition towards a clean energy system²⁴⁹. They have contended that the complex web of technologies, infrastructures, organizations, markets, regulations, and user practices required to deliver energy, heat, and transport services to society, must all transition for any meaningful reduction in global emissions. From the STT's "multi-level perspective" (MLP), any major socio-technical transition will be the outcome of three mutually reinforcing processes: (1) "Increasing momentum of niche innovations, (2) weakening of existing systems; and (2) strengthening exogenous pressures"²⁵⁰. All of these are required for a successful transition to a profoundly new technology. This has typically taken over a century for other major technological transitions, e.g. from sail to steam shipping²⁵¹. Accordingly, they have argued that accelerating technological progress requires not just greater investment in research and development, but enormous political support, widespread market and social acceptance, and a weakening of the existing incumbent regime. Such an approach provides a plausible narrative for the challenges facing the world in orchestrating a transition of the pace implied by the Fast Transition.

There is no disputing the vital role the global community must play in ensuring the deployment of clean technologies necessary to drive them down their experience curves. This is in fact the main motivation for this research. However, it is not the intention of this analysis to examine all such social and political barriers to the clean energy transition. Our intent is to examine what is possible and provide a proof of principle regarding the engineering challenge associated with the transition. We acknowledge there will be many such social, institutional, and material hurdles to overcome in a Fast Transition. Such hurdles could be listed and even modelled but what often proves more difficult is to model mankind's ability to overcome such hurdles. For instance, in World War II when Japan captured Malaya they also captured the largest producer of rubber in

the world. Within 11 months all vehicles produced in the United States had synthetic rubber tires. One of the benefits of relying on experience curves, and their apparent persistence, is their empirical record embodies not only how such barriers have slowed down the pace of progress, but also the pace at which those barriers have been overcome or sidestepped.

The fact that the Fast Transition uses the same deployment growth rates that have persisted for the last few decades gives us confidence that such a pathway is possible. It is also important to note that the Fast Transition scenario does not require the energy system to be built from scratch, but rather only the phasing in of new technologies into an existing infrastructure, a transition that has traditionally occurred much faster²⁵². Furthermore, if we can achieve the increased rate of electrification required for the Fast Transition, as discussed in S.I. 1.6.2, then the social and political challenges associated with the clean energy transition become considerably less daunting due to the relative ease in which electricity systems have transitioned to clean energy sources in the past²⁵³. Unlike the food, buildings, and transport sectors, changing electricity sources does not require significant consumer involvement. Electricity is an undifferentiated product, which means that consumers do not experience any change to their services whether the electrons powering their house are emissions-intensive or not. Policy makers have also found it easier to deal with decarbonising the power sector due to the relatively small number of utilities involved, in contrast to the millions of small farmers/owners/builders in the food, transport, and buildings sectors²⁵³. In fact, applying the MLP analysis presented by Geels et al.²⁵³ would suggest that for some early adopters like Germany and the UK the uptake of renewables into the electricity system is already past most of the barriers that might prevent the widespread adoption of such new technologies, including comparatively high initial technology costs, the lack of dominant designs, and public, corporate, or political opposition. Solar and wind both have decades of development behind them, so these technologies are now arguably in the final phase of their transition to mainstream energy supply. As per the multi-level perspective the final phase in these formerly-niche technologies becoming part of the mainstream regime will involve further adjustments in existing infrastructure (e.g. more storage), market structures (e.g. incumbent fight-back), and public perceptions. These processes are well underway in many countries and given price parity with existing fossil fuel technologies is likely in the coming few years for major economies such as India, China and the US (S.I. 8.6) would be difficult to stop.

This research is intended to add to the current body of evidence concerning the cost of the transition by applying a rigorous probabilistic technology forecast methodology to future technology costs that are not currently well represented in major climate mitigation models. Considerably more literature could be cited here expounding on the many social, institutional, and economic barriers to a fast, decisive clean energy transition, e.g. Nelson and Allwood²⁵⁴, Geels et al.²⁵⁰, Romero-Lankao et al.²³. The modelling approach used in this present study is much closer in theory to the STT perspective than that used by other approaches, such as cost optimisation models (S.I. 8.4). These latter models generally identify single optimal pathways that depend entirely on the precision of all the model's assumptions, and that often include technologies that are not yet feasible at scale, or are yet to gain social acceptance, such as bioenergy with carbon capture and storage²⁵³. In contrast, this present research relies on decade-long empirical deployment and cost trends for known technologies with persistent trends – trends that incorporate each technology's transition through niche markets, user acceptance, incumbent resistance, political wins and losses, and changes to infrastructure. This is the very substance of the narrative provided by the STT approach. The slope of the experience curve represents an aggregation of all such barriers being overcome in different markets, regions, and political environments, at different rates and times across the globe. The theoretical framing provided by the STT perspective actually provides a plausible narrative for the “random shocks” that hinder or accelerate progress of a technology, producing not only the shape of the experience curve but also the uncertainty in the empirical record that is used to generate the forecast errors we associate with them in this research.

A central underlying conjecture of this present research is that more attention should be paid to using empirically grounded approaches for technology forecasting in modelling energy transition pathways. There is clearly a balance between model realism, which results in increased model complexity, and model reliability, which demands lower complexity. Evidence suggests that large energy-economy models used to model transition pathways in the past have been, and perhaps still are, putting too much weight on aiming for realism, but in doing so are unable to capture basic technological trends – trends that may well end up rendering all their scenarios out-of-date.

8.7.1 Are the growth rates of electrolyzers in the Fast Transition plausible?

The wider research community has only recently begun to embrace the idea of hydrogen as a global energy carrier. There is therefore still considerable uncertainty surrounding the role that hydrogen might play in

the coming energy transition. Add to this the inherent complexity of the problem and we see a wide range of hydrogen technology deployment projections in future energy system scenarios²⁵⁵.

According to the IEA, electrolyser capacity will double from around 290MW in 2020 to nearly 700 MW of electrolyzer capacity commissioned in 2022¹¹⁰. The Hydrogen Council²⁵⁶ estimates the global electrolyzer capacity will continue to rise to around 12 GW by 2025, an estimate that aligns well with the Fast Transition. Using data from the IEA Hydrogen Technology Collaboration Programme, IRENA²⁰³ estimate that around 25 GW of water electrolysis projects will be operational by 2030, which is less than half of the cumulative deployment in the Fast Transition. However, this estimate is likely to be low given recent national announcements following COP26 including the EU's Hydrogen Strategy with its target of 40 GW of hydrogen electrolyzer capacity by 2030, the Chilean government's pledge of 25 GW by 2030, and India's new National Hydrogen Mission which alone could add around 22 GW of hydrogen production capacity by 2030²⁵⁷.

Looking into the more distant future the Hydrogen Council 2017(@ estimates that demand for hydrogen could increase to 80 EJ by 2050, a level that is reached in the Fast Transition in 2047. The amount of hydrogen production in the Fast Transition also aligns reasonably well with projections from a number of high-hydrogen-deployment model runs from the IPCC's AR5 database (Figure S77). The most ambitious hydrogen scenario in the IAM AR5 database is the MERGE_ETL_2011 - AMPERE2-450-NoCCS-OPT, which includes 138.5 EJ/yr of hydrogen production by 2050 (compared to the 118 EJ/yr deployed in the Fast Transition). Note that the IAM scenarios likely include hydrogen generation from many sources. For instance, the MERGE-ETL scenarios, which make up most of the high hydrogen model runs in AR5, involve a shift from initially generating mostly blue hydrogen to generating hydrogen using renewables, nuclear, and biofuels²⁵⁹.




Figure S77: Projected production of hydrogen in final energy (EJ/yr) from electrolysis in the Fast Transition compared with hydrogen production from all sources in a range of “high hydrogen” scenarios from the IPCC’s AR5 Database. Source for AR5 data: IAMC AR5 Scenario Database, 2014 available at: <https://secure.iiasa.ac.at/web-apps/ene/AR5DB/>

As discussed in S.I. 6.13, given the early stage of development of the key electrolyzer technology (PEMs) used for the Fast Transition cost projections, we apply a growth rate of around 60% from 2020 to 2040 for electrolytic hydrogen. If instead we applied the current approximately 100% growth rate for PEMs to electrolyzer deployment for one decade, followed by a much lower 8% growth rate from 2030 onwards we would achieve a similar cumulative production of hydrogen by 2070, and produce a reasonable match to the projections of hydrogen production in the IPCC IAM GRAPE v2011 EMF27-550-LimBio scenario (Figure S78).

The Fast Transition therefore appears to align reasonably well with existing high growth estimates for future hydrogen production. However, Blanco and Faaij²⁶⁰ question the credibility of electrolyzers sustaining the high annual deployment growth rates seen in solar PV, although they provide little reasoning. There is little doubt that such high deployment rates will require significant policy support. The economic tipping point of self-sustained growth, where renewable electrolytic hydrogen could become cheaper than blue or grey alternatives, is estimated to require an aggregated deployment of around 70 GW of electrolyzer capacity²⁵⁶. In our Fast Transition this would not occur until around 2029, suggesting less than a decade of support will be required to sustain the high growth rates in the Fast Transition before reaching price parity.

Developing sufficient capacity to deploy the quantities of hydrogen technology in the Fast Transition will also require us to develop off-taker markets and full-featured infrastructure systems that can service and supply P2X fuels for heavy-duty vehicles, ships, and applications in industries, such as steel and cement, and building heating²⁶¹. However, the faster such a roll-out occurs the faster unit investment costs required to

switch to a green-fuel-supported economy may decline with hydrogen production becoming streamlined and transmission and distribution infrastructure utilised more effectively²⁶¹. Introducing hydrogen into existing gas grids is already underway (e.g. South Australia), which might enable a smoother transition from existing markets. Large-scale hydrogen storage options for the various sectors should also be explored to reduce uncertainty in the different global markets.




Figure S78: Cumulative hydrogen production (EJ/yr) for the Fast Transition compared with an alternative growth path consisting of a 100% deployment growth rate from 2020 to 2030 followed by an 8% deployment growth rate from 2030 to 2070, which aligns well with the projections of the IPCC AR5 GRAPE v2011 EMF27-550-LimBio scenario. Source for GRAPE data points: IAMC AR5 Scenario Database, 2014 available at: <https://secure.iiasa.ac.at/web-apps/ene/AR5DB/>

8.8 Additional benefits from the Fast Transition

8.8.1 Emission reductions and reduced climate risks

One of the clearest benefits of the Fast Transition is the rapid decline in emissions it provides, as shown in Figure S79a. This scenario achieves close to zero emissions by 2050 for all of the energy sector components represented in the model, contrasting strongly with the protracted reductions in the Slow Transition and the steady rise in the No Transition. Reducing emissions rapidly in the first decade is important to limit the cumulative emissions from pre-industrial levels and associated warming by 2070, as shown with the dark bars in Figure S79b.

To put these differences in terms of global warming scenarios requires accounting for all anthropogenic emissions. The model accounts for around 82% of the CO₂ emissions from the energy supplied to all end-use sectors, which amounts to around 30 gigatonnes (Gt) of the approximately 50Gt of total global anthropogenic emissions in 2020. The energy-system components not included (e.g. oil for electricity, bioenergy for heating and cooking, methane leakage, etc.) amount to around 8.1Gt of emissions in 2020. Non-energy emissions in 2020 add a further 10Gt of GHG emissions, including cement and steel production and agriculture, forestry, and other land use changes (AFOLU). Producing a rough estimate of the warming associated with each of our scenarios requires a comparison of our model with that of a full global energy system model, such as provided by the International Energy Agency, as well as assumptions around the fate of non-energy emissions. This exercise is documented in detail in Ives et al.²⁶², and the estimated average additional cumulative emissions for each scenario are presented in Figure S79b in the white boxes above each scenario's modelled cumulative

emissions bar. The estimated warming in 2070 associated with all of these cumulative emission estimates (calculated using FaIR^{263,264}) are shown as temperature anomalies using the axis of Figure S79c, which also presents the temperature anomalies associated with several IPCC Representative Concentration Pathways in 2070 for comparison.

To summarise the key message of the evidence presented in Figure S79, the Fast Transition enables the modelled global energy system to achieve an emissions reduction pathway consistent with the Paris Agreement of 1.5 degrees of warming by 2100 (RCP1.9 - light blue shaded area in Figure S79c). However, without similar strong efforts to decarbonise the remaining sources of anthropogenic emissions from the global economy, such as land use change, cement production, and industrial processes, global warming could still exceed 2 degrees above pre-industrial levels by 2100. Achieving a Paris-compliant pathway in the Slow Transition will require even greater efforts in non-energy sectors, including applying potentially expensive negative emissions technologies.




Figure S79: Comparison of emissions from modelled scenarios including a) a comparisons of estimated global annual CO₂ emissions from 2020 to 2070 for the modelled No Transition (black), Slow Transition (blue), and Fast Transition (green) scenarios; b) the cumulative emissions from pre-industrial levels for each of the scenarios including the modelled components of the global energy system (dark bar areas) and a mean estimate of the additional greenhouse gas emissions associated with components of the global economic system not included in our energy system model (white bar areas) and c) the range of modelled temperature anomalies associated with greenhouse gas emissions for the IPCC Representative Concentration Pathways 1.9, 2.6, 4.5, 6.0 and 8.5 in 2070. Comparing the full bars for each scenario with the RCP shaded bands provides a rough guide as to which set of RCP scenarios is most closely aligned with each scenario in 2070. The RCP1.9 scenario (aqua blue) is made up from modelled pathways designed to meet the Paris Agreement’s “well below 2 degrees” (with an ambition for 1.5 degrees) of warming above pre-industrial levels by 2100. The mean cumulative additional emissions, shown as white boxes in b), are generated using a range of alternative assumptions around the components not included in this model, as detailed in Ives et al.²⁶².

The benefits of rapidly reducing greenhouse gas emissions and averting the worst impacts of climate change are the subject of increasing research and debate. What is evident is that the world is already experiencing more extreme weather events and adverse impacts attributable to warming temperatures^{265,266}.

Attempts to quantify the costs associated with potential future climate damages have focused on an evaluation of what is termed the social cost of carbon (SCC), which represents the present social value of damages from emitting CO₂, and the associated warming anomalies it generates. Unfortunately, our understanding of how climate change might impact on societies and economies is far less developed than our understanding of climate systems²⁶⁷. Estimates of the SCC have been controversial from the very beginning. In 2009, the

Obama Administration tasked an Inter-agency Working Group (IWG) with developing a robust SCC, based on the best available science and economics. In 2010 the IWG²⁶⁸ produced an estimated value of around \$53 per ton of CO₂ in 2020. However, significant disparity exists between such standard economic damage estimates, which conclude relatively minor impairments to the global economy, and more catastrophic projections, particularly those from the non-traditional economic and climate sciences^{269,270}. The possibility of catastrophic damages at higher temperatures could necessitate modelling impacts using non-linear damage functions, which are likely to result in much greater potential losses²⁷¹. For example, incorporating the possibility of non-linear tipping points in the climate system has the estimated potential to increase SCC estimates by a factor of eight²⁷². A more recent study, drawing on the latest climate projections and more robust country-level macroeconomic modelling calculated a median estimate for the global SCC in 2018 at around \$417 per ton²⁷³.

In providing an estimate of the potential costs associated with climate change we have again taken a relatively conservative approach. We evaluate the climate damages associated with only the emissions associated with the energy system components included in our model, and use a SCC range of 30 - 300 \$/tCO₂ (rising at 3% per year²⁷⁴). Calculated climate damage estimates are provided for each scenario at a 1.4% and 5% discount rate, in Table S27, along with expected energy system costs, and total combined costs. Subtracting the Fast Transition from the No Transition estimates at a discount rate of 5%, yields total expected Fast Transition savings, up to 2070, of 31 - 255 trillion dollars. At the lower discount rate of 1.4%, the range of expected savings is 88 - 775 trillion dollars. Thus the benefits of the Fast Transition are likely to be much greater than the energy system cost savings we have evaluated in this study.

Scenario	Discount rate (%)	SCC (\$/tCO ₂)	Climate damages (\$tn)	Expected NPC of modelled energy system (\$tn)	Total (\$tn)
Fast Transition	1.4	30	20	195	215
Fast Transition	1.4	300	201	195	396
Slow Transition	1.4	30	48	202	250
Slow Transition	1.4	300	478	202	680
No Transition	1.4	30	96	207	303
No Transition	1.4	300	964	207	1171
Fast Transition	5.0	30	14	88	102
Fast Transition	5.0	300	139	88	227
Slow Transition	5.0	30	25	92	117
Slow Transition	5.0	300	246	92	338
No Transition	5.0	30	39	94	133
No Transition	5.0	300	389	94	482

Table S27: Summary of climate damage estimates up to 2070 for each of the main scenarios across a range of discount rates and SCC values. (Dollars are constant 2020 dollars, costs are rounded to nearest integer.)

8.8.2 Co-benefits of a Fast Transition

There are also other important benefits of accelerating the green transition beyond those directly associated with reducing adverse climate change impacts²⁷⁵. The most obvious co-benefit is reduced mortality and morbidity from the elimination of the air pollution associated with fossil fuel use. For example, the World Health Organization estimates that 4.2 million lives per year are lost to ambient air pollution, caused mostly by fossil fuels²⁷⁶.

The volatility of fossil fuel prices also has significant adverse impacts on the global economy. The renewable energy transition will likely substantially reduce the volatility of energy prices²⁷⁷. From a geopolitical point of view, the use of renewables together with P2X fuels might also offer energy security to many countries that are currently dependent on others. Finally, providing cheap renewable energy to low income countries can accelerate energy inclusion and contribute to meeting many of the Sustainable Development Goals, including lifting the estimated 1 billion people out of energy poverty²³.

Finally, a rapid decarbonisation of the energy system dramatically reduces our annual emissions, potentially buying more time to develop approaches for decarbonising the much harder-to-abate sectors, such as agriculture, steel, cement, and air transport.

8.9 Additional costs from the Fast Transition

8.9.1 Stranded Assets

A stranded asset has been defined as “an asset which loses economic value well ahead of its anticipated useful life, whether that is a result of changes in legislation, regulation, market forces, disruptive innovation, societal norms, or environmental shocks”²⁷⁸. However, different disciplines essentially have different definitions based on their main area of concern, with accountants looking for “impairment”; economists looking for whether the social benefits exceed economic costs of maintaining an asset; regulators looking for “stranded costs”; and investors being most concerned with financial losses²⁷⁹. All such definitions could arrive at a different point in the engineering lifetime of an asset when they would consider the asset stranded, particularly when repairs and maintenance can prolong that asset’s lifetime.

Tong et al.²⁸⁰ provide a comprehensive assessment of the potential costs of stranded assets from meeting a 1.5 degrees target. This analysis suggests that global emissions could exceed the 1.5 degrees carbon budget (with a probability of 66 to 50%) if the more than 500 GW of fossil-fuel based power and transport assets currently in planning or proposed were built and operated for their full engineering lifetimes. They calculate that a scenario that meets the 1.5 degrees target will likely involve a stranding of \$5 to \$17 trillion in asset value. The gap between the low and high estimates is mostly composed of stranded assets in the transport sector, which have high value losses per tonne of emissions reductions (or the lowest committed emissions per unit asset value), predominantly through fossil fuel-based passenger vehicles becoming stranded. The highest value is therefore from a scenario where transport assets are stranded rather than the assets with high committed emissions per unit asset, like coal-fired power stations.

This highlights one key issue with such estimates of the value of stranded assets. Unlike energy producing technologies most passenger cars do not generate income, but rather provide utility through the provision of a service e.g. mobility. This makes an evaluation of their stranded value more difficult to calculate as their market value, or people’s willingness to pay for the service they provide, can drop dramatically when a new superior technology becomes available, such as electric vehicles or a subscription-based robot-taxi service. If that alternative technology can provide the same or better service at a lower cost, the value of the original asset will experience a natural decline in its asset market value.

Another issue with such calculations of stranded assets is the inclusion of planned capacity. The work of Tong et al.²⁸⁰ is an update on previous work, such as that published in the previous year by Pfeiffer et al.²⁸¹. In comparing these two works we can see a marked (30%) decline in their estimates of proposed/planned fossil fuel investments. According to Global Energy Monitor this is the continuation of a potentially growing trend from 2010, which has seen only 35% of planned capacity built or starting construction in 2019 (993GW), with some 1815GW of capacity cancelled or shelved²⁸². The more such trends reduce the 500 GW of planned and proposed assets included in the Tong et al.²⁸⁰ estimates the less likely there will be stranded assets.

Finally, it is important to differentiate between asset that are stranded due to policies that stop polluting activities, and assets stranding due to technological obsolescence, as the two are qualitatively different. Obsolescence is a natural outcome of creative destruction in capitalist economies – the incessant innovation mechanism by which new products and processes replace outdated ones²⁸³. Many products, such as computers, monitors, phones, printers, light bulbs, etc. are continually being replaced by superior alternatives before their engineering lifetimes.

This is particularly the case in the faster transition scenarios where the costs of renewables decline to the point where polluting assets are initially underutilised, and then replaced early as the cost to replace them with renewable sources is less than the fuel, operation and maintenance costs of the existing asset. Tong et al.²⁸⁰ assume historic rates of utilisation for such assets in future, which would overestimate their emissions, particularly since electricity generation markets are becoming increasingly competitive. Thus, because the Fast Transition involves moving to cheaper, cleaner and increasingly more popular sources of energy, and the Paris Agreement was ratified by all countries in 2016, most of the early closures could be regarded as the result of sub-optimal investment decisions, rather than forced closures.

Determining which of the existing assets can be truly regarded as “stranded” from a market perspective or even a social planning perspective – taking into account natural product obsolescence, and imperfect investment decision making – is beyond the scope of this paper. However, there is sufficient evidence to suggest that the true value of stranded assets will be less than the estimates provided in Tong et al.²⁸⁰, and certainly orders of magnitude less than even a conservative estimate of the climate damages associated with a No Transition or Historical Mix scenario²⁶².

References

- [1] Grubb, M., Drummond, P., Poncia, A., McDowall, W., Popp, D., Samadi, S., Penasco, C., Gillingham, K. T., Smulders, S., Glachant, M., Hassall, G., Mizuno, E., Rubin, E. S., Dechezleprêtre, A., and Pavan, G. (2021). Induced innovation in energy technologies and systems: a review of evidence and potential implications for CO₂ mitigation. *Environmental Research Letters*, 16(4):043007. doi: 10.1088/1748-9326/abde07.
- [2] Lafond, F., Bailey, A. G., Bakker, J. D., Rebois, D., Zadourian, R., McSharry, P., and Farmer, J. D. (2018). How well do experience curves predict technological progress? a method for making distributional forecasts. *Technological Forecasting and Social Change*, 128:104–117.
- [3] Gritsevskiy, A. and Nakicenovic, N. (2000). Modeling uncertainty of induced technological change. *Energy Policy*, 28(13):907–17909.
- [4] Gambhir, A., Butnar, I., Li, P.-H., Smith, P., and Strachan, N. (2019). A review of criticisms of integrated assessment models and proposed approaches to address these, through the lens of beccs. *Energies*, 12(9). doi: 10.3390/en12091747.
- [5] Pindyck, R. S. (2013). Climate change policy: What do the models tell us? *Journal of Economic Literature*, 51(3):860–72. doi: 10.1257/jel.51.3.860.
- [6] Farmer, J. D., Hepburn, C., Mealy, P., and Teytelboym, A. (2015). A third wave in the economics of climate change. *Environmental and Resource Economics*, 62(2):329–357.
- [7] Stanton, E. A., Ackerman, F., and Kartha, S. (2009). Inside the integrated assessment models: Four issues in climate economics. *Climate and Development*, 1(2):166–184. doi: 10.3763/cdev.2009.0015.
- [8] GEA (2012). *Global Energy Assessment - Toward a Sustainable Future*. Cambridge University Press and the International Institute for Applied Systems Analysis.
- [9] Wilson, C., Grubler, A., Gallagher, K. S., and Nemet, G. F. (2012). Marginalization of end-use technologies in energy innovation for climate protection. *Nature Climate Change*, 2(11):780–788. doi: 10.1038/nclimate1576.
- [10] Creutzig, F., Roy, J., Lamb, W. F., Azevedo, I. M. L., Bruine de Bruin, W., Dalkmann, H., Edelenbosch, O. Y., Geels, F. W., Grubler, A., Hepburn, C., Hertwich, E. G., Khosla, R., Mattauch, L., Minx, J. C., Ramakrishnan, A., Rao, N. D., Steinberger, J. K., Tavoni, M., Ürge-Vorsatz, D., and Weber, E. U. (2018). Towards demand-side solutions for mitigating climate change. *Nature Climate Change*, 8(4):260–263. doi: 10.1038/s41558-018-0121-1.
- [11] Lovins, A. B. (2018). How big is the energy efficiency resource? *Environmental Research Letters*, 13(9):090401. doi: 10.1088/1748-9326/aad965.
- [12] IEA (2021). *World Energy Outlook 2021*. Technical report, International Energy Agency.
- [13] IEA (2019). *World Energy Outlook 2019*. Technical report, International Energy Agency.
- [14] UN FAO (2009). *The State of food and agriculture, 2008*. Technical Report 11, Food and Agriculture Organization of the United Nations. URL <http://www.fao.org/3/i0100e/i0100e.pdf>.
- [15] Bui, M., Adjiman, C. S., Bardow, A., Anthony, E. J., Boston, A., Brown, S., Fennell, P. S., Fuss, S., Galindo, A., Hackett, L. A., Hallett, J. P., Herzog, H. J., Jackson, G., Kemper, J., Krevor, S., Maitland, G. C., Matuszewski, M., Metcalfe, I. S., Petit, C., Puxty, G., Reimer, J., Reiner, D. M., Rubin, E. S., Scott, S. A., Shah, N., Smit, B., Trusler, J. P. M., Webley, P., Wilcox, J., and Mac Dowell, N. (2018). Carbon capture and storage (ccs): the way forward. *Energy Environ. Sci.*, 11:1062–1176. doi: 10.1039/C7EE02342A.
- [16] Hill, B., Hovorka, S., and Melzer, S. (2013). Geologic carbon storage through enhanced oil recovery. *Energy Procedia*, 37:6808–6830. doi: <https://doi.org/10.1016/j.egypro.2013.06.614>.
- [17] Rubin, E. S., Davison, J. E., and Herzog, H. J. (2015). The cost of CO₂ capture and storage. *International Journal of Greenhouse Gas Control*, 40:378 – 400. doi: <https://doi.org/10.1016/j.ijggc.2015.05.018>.

- [18] Budinis, S., Krevor, S., Dowell, N. M., Brandon, N., and Hawkes, A. (2018). An assessment of ccs costs, barriers and potential. *Energy Strategy Reviews*, 22:61–81. doi: <https://doi.org/10.1016/j.esr.2018.08.003>.
- [19] Reiner, D. M. (2016). Learning through a portfolio of carbon capture and storage demonstration projects. *Nature Energy*, 1(1):1–7.
- [20] Ricardo (2020). Analysing the potential of bioenergy with carbon capture in the UK to 2050. Technical report, Ricardo Energy & Environment.
- [21] CCC (2019). Net Zero: The UK’s contribution to stopping global warming. Technical report, Committee on Climate Change.
- [22] Sachs, J. D. (2015). What’s The Path to Deep Decarbonization? URL <https://www.weforum.org/agenda/2015/12/whats-the-path-to-deep-decarbonization/>.
- [23] Romero-Lankao, P., Wilson, A., Sperling, J., Miller, C., Zimny-Schmitt, D., Sovacool, B., Gearhart, C., Muratori, M., Bazilian, M., Zünd, D., Young, S., Brown, M., and Arent, D. (2021). Of actors, cities and energy systems: advancing the transformative potential of urban electrification. *Progress in Energy*, 3(3):032002. doi: 10.1088/2516-1083/abfa25.
- [24] IEA (2019). Key World Energy Statistics. Technical report, International Energy Agency.
- [25] Brynolf, S., Taljegard, M., Grahn, M., and Hansson, J. (2018). Electrofuels for the transport sector: A review of production costs. *Renewable and Sustainable Energy Reviews*, 81:1887 – 1905. doi: <https://doi.org/10.1016/j.rser.2017.05.288>.
- [26] IEA (2019). The Future of Hydrogen. Technical report, International Energy Agency.
- [27] WEC (2018). Bringing north sea energy ashore efficiently. Technical report, World Energy Council.
- [28] Hydrogen Council, and McKinsey and Company (2021). Hydrogen Insights: A perspective on hydrogen investment, market development and cost competitiveness. Technical report.
- [29] BNEF (2020). Electric Vehicle Outlook 2020.
- [30] IEA (2021). Global EV Outlook 2021 - Accelerating ambitions despite the pandemic.
- [31] Nelder, C. and Rogers, E. (2019). Reducing EV Charging Infrastructure Costs. Technical report, Rocky Mountain Institute.
- [32] IRS (2021). Retail audit technique guide. Technique Guide Publication 5495 (3-2021), Internal Revenue Service, Department of the Treasury.
- [33] Arthur, W. B. (2011). *The Nature of Technology. What It Is and How It Evolves*. Free Press.
- [34] Schlömer, S., Bruckner, T., Fulton, L., Hertwich, E., McKinnon, A., Perczyk, D., Roy, J., Schaeffer, R., Sims, R., Smith, P., and Wisser, R. Annex iii: Technology-specific cost and performance parameters. In Edenhofer, O., Pichs-Madruga, R., Sokona, Y., Farahani, E., Kadner, S., Seyboth, K., Adler, A., Baum, I., Brunner, S., Eickemeier, P., Kriemann, B., Savolainen, J., Schlömer, S., von Stechow, C., Zwickel, T., and Minx, J., editors, *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2014.
- [35] Barbose, G., Darghouth, N., Elmallah, S., Forrester, S., LaCommare, K., Millstein, D., Rand, J., Cotton, W., Sherwood, S., and O’Shaughnessy, E. (2019). Tracking the sun 2019, pricing and design trends for distributed photovoltaic systems in the united states 2019 edition. Technical report, Lawrence Berkeley National Laboratory.
- [36] Barbose, G., Darghouth, N., O’Shaughnessy, E., and Forrester, S. (2021). Tracking the sun 2021, pricing and design trends for distributed photovoltaic systems in the united states 2021 edition. Technical report, Lawrence Berkeley National Laboratory.

- [37] Brown, T., Schlachtberger, D., Kies, A., Schramm, S., and Greiner, M. (2018). Synergies of sector coupling and transmission reinforcement in a cost-optimised, highly renewable european energy system. *Energy*, 160:720–739. doi: <https://doi.org/10.1016/j.energy.2018.06.222>.
- [38] Brown, T. and Reichenberg, L. (2021). Decreasing market value of variable renewables can be avoided by policy action. *Energy Economics*, 100:105354. doi: <https://doi.org/10.1016/j.eneco.2021.105354>.
- [39] OECD & NEA (2019). The costs of decarbonisation: System costs with high shares of nuclear and renewables. Technical report, OECD and Nuclear Energy Agency.
- [40] Trancik, J. E., Jean, J., Kavlak, G., Klemun, M. M., Edwards, M. R., McNerney, J., Miotti, M., Brown, P. R., Mueller, J. M., and Needell, Z. A. (2015). Technology improvement and emissions reductions as mutually reinforcing efforts: Observations from the global development of solar and wind energy. Working paper, MIT.
- [41] De Stercke, S. (2014). Dynamics of energy systems: A useful perspective. IIASA interim report, IIASA. URL <http://pure.iiasa.ac.at/id/eprint/11254/>.
- [42] Staffell, I., Brett, D., Brandon, N., and Hawkes, A. (2012). A review of domestic heat pumps. *Energy Environ. Sci.*, 5:9291–9306. doi: 10.1039/C2EE22653G.
- [43] Madeddu, S., Ueckerdt, F., Pehl, M., Peterseim, J., Lord, M., Kumar, K. A., Krüger, C., and Luderer, G. (2020). The CO₂reduction potential for the European industry via direct electrification of heat supply (power-to-heat). *Environmental Research Letters*, 15(12):124004. doi: 10.1088/1748-9326/abbd02.
- [44] Noel, L., Zarazua de Rubens, G., Kester, J., and Sovacool, B. E. (2019). *Vehicle-to-Grid: A Sociotechnical Transition Beyond Electric Mobility*. Basingstoke: Palgrave MacMillan.
- [45] Lovins, A. B. (2017). Reliably integrating variable renewables: Moving grid flexibility resources from models to results. *The Electricity Journal*, 30(10):58 – 63. doi: <https://doi.org/10.1016/j.tej.2017.11.006>.
- [46] Grubler, A., Wilson, C., Bento, N., Boza-Kiss, B., Krey, V., McCollum, D. L., Rao, N. D., Riahi, K., Rogelj, J., De Stercke, S., Cullen, J., Frank, S., Fricko, O., Guo, F., Gidden, M., Havlík, P., Huppmann, D., Kiesewetter, G., Rafaj, P., Schoepp, W., and Valin, H. (2018). A low energy demand scenario for meeting the 1.5 °c target and sustainable development goals without negative emission technologies. *Nature Energy*, 3(6):515–527. doi: 10.1038/s41560-018-0172-6.
- [47] Grubler, A., Nakićenović, N., and Victor, D. G. (1999). Dynamics of energy technologies and global change. *Energy Policy*, 27(5):247–280. doi: 10.1016/S0301-4215(98)00067-6.
- [48] Ritchie, H. and Roser, M. (2018). Our world in data, energy production and changing energy sources. URL <https://ourworldindata.org/energy-production-and-changing-energy-sources>. [Online; accessed 22-December-2019].
- [49] Mai, T., Jadun, P., Logan, J., McMillan, C., Muratori, M., Steinberg, D., Vimmerstedt, L., Jones, R., Haley, B., and Nelson, B. (2018). *Electrification futures study: Scenarios of electric technology adoption and power consumption for the united states*. Nrel/tp-6a20-71500, NREL. URL <https://www.nrel.gov/docs/fy18osti/71500.pdf>.
- [50] Kersey, J., Popovich, N. D., and Phadke, A. A. (2022). Rapid battery cost declines accelerate the prospects of all-electric interregional container shipping. *Nature Energy*, 7(7):664–674. doi: 10.1038/s41560-022-01065-y.
- [51] Deason, J., Wei, M., Leventis, G., Smith, S., and Schwartz, L. C. (2018). *Electrification of buildings and industry in the united states: Drivers, barriers, prospects, and policy approaches*. Technical report, LBNL.
- [52] Kassakian, J. G., Schmalensee, R., Hogan, W. W., Jacoby, H. D., and Kirtley, J. L. (2011). *The future of the electric grid*. Technical report, MIT.
- [53] IEA (2010). *Power Generation from Coal - Measuring and Reporting Efficiency Performance and CO₂ Emissions*. Technical report, International Energy Agency.

- [54] IEA (2015). Coal-fired power plant efficiency improvement in India. Technical report, International Energy Agency.
- [55] Pedersen, T. T., Victoria, M., Rasmussen, M. G., and Andresen, G. B. (2021). Modeling all alternative solutions for highly renewable energy systems. *Energy*, 234:121294. doi: <https://doi.org/10.1016/j.energy.2021.121294>.
- [56] Neumann, F. and Brown, T. (2021). The near-optimal feasible space of a renewable power system model. *Electric Power Systems Research*, 190:106690. doi: <https://doi.org/10.1016/j.eprsr.2020.106690>.
- [57] Blanco, H. and Faaij, A. (2018). A review at the role of storage in energy systems with a focus on power to gas and long-term storage. *Renewable and Sustainable Energy Reviews*, 81:1049 – 1086. doi: <https://doi.org/10.1016/j.rser.2017.07.062>.
- [58] Sepulveda, N. A., Jenkins, J. D., de Sisternes, F. J., and Lester, R. K. (2018). The role of firm low-carbon electricity resources in deep decarbonization of power generation. *Joule*, 2(11):2403–2420. doi: <https://doi.org/10.1016/j.joule.2018.08.006>.
- [59] Bistline, J., Cole, W., Damato, G., DeCarolis, J., Frazier, W., Linga, V., Marcy, C., Namovicz, C., Podkaminer, K., Sims, R., Sukunta, M., and Young, D. (2020). Energy storage in long-term system models: a review of considerations, best practices, and research needs. *Progress in Energy*, 2(3):032001. doi: 10.1088/2516-1083/ab9894.
- [60] McCollum, D. L., Gambhir, A., Rogelj, J., and Wilson, C. (2020). Energy modellers should explore extremes more systematically in scenarios. *Nature Energy*, 5(2):104–107. doi: 10.1038/s41560-020-0555-3.
- [61] Sepulveda, N. A., Jenkins, J. D., Edington, A., Mallapragada, D. S., and Lester, R. K. (2021). The design space for long-duration energy storage in decarbonized power systems. *Nature Energy*, 6(5):506–516. doi: 10.1038/s41560-021-00796-8.
- [62] IRENA (2022). Renewable Power Generation Costs in 2021. Technical report, International Renewable Energy Agency.
- [63] Victoria, M., Zhu, K., Brown, T., Andresen, G. B., and Greiner, M. (2019). The role of storage technologies throughout the decarbonisation of the sector-coupled european energy system. *Energy Conversion and Management*, 201:111977. doi: <https://doi.org/10.1016/j.enconman.2019.111977>.
- [64] Gea-Bermúdez, J., Jensen, I. G., Münster, M., Koivisto, M., Kirkerud, J. G., Kuang Chen, Y., and Ravn, H. (2021). The role of sector coupling in the green transition: A least-cost energy system development in northern-central europe towards 2050. *Applied Energy*, 289:116685. doi: <https://doi.org/10.1016/j.apenergy.2021.116685>.
- [65] Kittner, N., Castellanos, S., Hidalgo-Gonzalez, P., Kammen, D. M., and Kurtz, S. (2021). Cross-sector storage and modeling needed for deep decarbonization. *Joule*, 5(10):2529–2534. doi: <https://doi.org/10.1016/j.joule.2021.09.003>.
- [66] Ziegler, M. S., Mueller, J. M., Pereira, G. D., Song, J., Ferrara, M., Chiang, Y.-M., and Trancik, J. E. (2019). Storage requirements and costs of shaping renewable energy toward grid decarbonization. *Joule*, 3(9): 2134–2153. doi: <https://doi.org/10.1016/j.joule.2019.06.012>.
- [67] Dowling, J. A., Rinaldi, K. Z., Ruggles, T. H., Davis, S. J., Yuan, M., Tong, F., Lewis, N. S., and Caldeira, K. (2020). Role of long-duration energy storage in variable renewable electricity systems. *Joule*, 4(9): 1907–1928. doi: <https://doi.org/10.1016/j.joule.2020.07.007>.
- [68] Albertus, P., Manser, J. S., and Litzelman, S. (2020). Long-duration electricity storage applications, economics, and technologies. *Joule*, 4(1):21–32. doi: <https://doi.org/10.1016/j.joule.2019.11.009>.
- [69] Guerra, O. J., Eichman, J., and Denholm, P. (2021). Optimal energy storage portfolio for high and ultrahigh carbon-free and renewable power systems. *Energy Environ. Sci.*, 14:5132–5146. doi: 10.1039/D1EE01835C.

- [70] Hunter, C. A., Penev, M. M., Reznicek, E. P., Eichman, J., Rustagi, N., and Baldwin, S. F. (2021). Techno-economic analysis of long-duration energy storage and flexible power generation technologies to support high-variable renewable energy grids. *Joule*, 5(8):2077–2101. doi: <https://doi.org/10.1016/j.joule.2021.06.018>.
- [71] Shaner, M. R., Davis, S. J., Lewis, N. S., and Caldeira, K. (2018). Geophysical constraints on the reliability of solar and wind power in the united states. *Energy Environ. Sci.*, 11:914–925. doi: 10.1039/C7EE03029K.
- [72] Tong, D., Farnham, D. J., Duan, L., Zhang, Q., Lewis, N. S., Caldeira, K., and Davis, S. J. (2021). Geophysical constraints on the reliability of solar and wind power worldwide. *Nature Communications*, 12(1):6146. doi: 10.1038/s41467-021-26355-z.
- [73] Grams, C. M., Beerli, R., Pfenninger, S., Staffell, I., and Wernli, H. (2017). Balancing europe’s wind-power output through spatial deployment informed by weather regimes. *Nature Climate Change*, 7(8): 557–562. doi: 10.1038/nclimate3338.
- [74] Zeyringer, M., Price, J., Fais, B., Li, P.-H., and Sharp, E. (2018). Designing low-carbon power systems for great britain in 2050 that are robust to the spatiotemporal and inter-annual variability of weather. *Nature Energy*, 3(5):395–403. doi: 10.1038/s41560-018-0128-x.
- [75] Jerez, S., Tobin, I., Turco, M., Jiménez-Guerrero, P., Vautard, R., and Montávez, J. (2019). Future changes, or lack thereof, in the temporal variability of the combined wind-plus-solar power production in europe. *Renewable Energy*, 139:251–260. doi: <https://doi.org/10.1016/j.renene.2019.02.060>.
- [76] Tröndle, T., Lilliestam, J., Marelli, S., and Pfenninger, S. (2020). Trade-offs between geographic scale, cost, and infrastructure requirements for fully renewable electricity in europe. *Joule*, 4(9):1929–1948. doi: <https://doi.org/10.1016/j.joule.2020.07.018>.
- [77] Brouwer, A. S., van den Broek, M., Zappa, W., Turkenburg, W. C., and Faaij, A. (2016). Least-cost options for integrating intermittent renewables in low-carbon power systems. *Applied Energy*, 161:48 – 74. doi: <https://doi.org/10.1016/j.apenergy.2015.09.090>.
- [78] Denholm, P. and Mai, T. (2019). Timescales of energy storage needed for reducing renewable energy curtailment. *Renewable Energy*, 130:388 – 399. doi: <https://doi.org/10.1016/j.renene.2018.06.079>.
- [79] Solomon, A., Kammen, D. M., and Callaway, D. (2014). The role of large-scale energy storage design and dispatch in the power grid: A study of very high grid penetration of variable renewable resources. *Applied Energy*, 134:75 – 89. doi: <https://doi.org/10.1016/j.apenergy.2014.07.095>.
- [80] Saarinen, L., Dahlbäck, N., and Lundin, U. (2015). Power system flexibility need induced by wind and solar power intermittency on time scales of 1–14 days. *Renewable Energy*, 83:339 – 344. doi: <https://doi.org/10.1016/j.renene.2015.04.048>.
- [81] Ziegler, M. S., Mueller, J. M., Pereira, G. D., Song, J., Ferrara, M., Chiang, Y.-M., and Trancik, J. E. (2019). Storage requirements and costs of shaping renewable energy toward grid decarbonization. *Joule*, 3(9): 2134 – 2153. doi: <https://doi.org/10.1016/j.joule.2019.06.012>.
- [82] Ziegler, M. S. and Trancik, J. E. (2021). Re-examining rates of lithium-ion battery technology improvement and cost decline. *Energy Environ. Sci.*, 14:1635–1651. doi: 10.1039/D0EE02681F.
- [83] Cole, W., Frazier, A. W., and Augustine, C. (2021). Cost Projections for Utility- Scale Battery Storage: 2021 Update. Technical report, NREL.
- [84] Ha, S. and Gallagher, K. G. (2015). Estimating the system price of redox flow batteries for grid storage. *Journal of Power Sources*, 296:122 – 132. doi: <https://doi.org/10.1016/j.jpowsour.2015.07.004>.
- [85] Li, Z., Pan, M. S., Su, L., Tsai, P.-C., Badel, A. F., Valle, J. M., Eiler, S. L., Xiang, K., Brushett, F. R., and Chiang, Y.-M. (2017). Air-breathing aqueous sulfur flow battery for ultralow-cost long-duration electrical storage. *Joule*, 1(2):306 – 327. doi: <https://doi.org/10.1016/j.joule.2017.08.007>.
- [86] EUTurbines (2019). EUTurbines renewable gas commitments 2019. URL <https://powertheeu.eu/>. [Online; accessed 22-December-2019].

- [87] Kobayashi, H., Hayakawa, A., Somarathne, K., and Okafor, E. (2019). Science and technology of ammonia combustion. *Proceedings of the Combustion Institute*, 37(1):109 – 133. doi: <https://doi.org/10.1016/j.proci.2018.09.029>.
- [88] EIA (2021). Annual Energy Outlook reports 1979-2021. Technical report, U.S. Energy Information Administration.
- [89] IEA (2021). Global Hydrogen Review 2021. Technical report, International Energy Agency.
- [90] Ikäheimo, J., Kiviluoma, J., Weiss, R., and Holttinen, H. (2018). Power-to-ammonia in future north european 100% renewable power and heat system. *International Journal of Hydrogen Energy*, 43(36): 17295 – 17308. doi: <https://doi.org/10.1016/j.ijhydene.2018.06.121>.
- [91] Icerman, L. (1974). Relative costs of energy transmission for hydrogen, natural gas, and electricity. *Energy Sources*, 1(4):435–446. doi: 10.1080/00908317408945936.
- [92] Brown, D., Cabe, J., and Stout, T. (2011). National lab uses ogj data to develop cost equations. *Oil and Gas Journal*, 109(1):108 – 111.
- [93] Ogden, J. M. (2018). Prospects for hydrogen in the future energy system. Technical report, Institute of Transportation Studies, University of California, Davis.
- [94] van der Zwaan, B., Schoots, K., Rivera-Tinoco, R., and Verbong, G. (2011). The cost of pipelining climate change mitigation: An overview of the economics of ch₄, co₂ and h₂ transportation. *Applied Energy*, 88(11):3821 – 3831. doi: <https://doi.org/10.1016/j.apenergy.2011.05.019>.
- [95] Dodds, P. E., Staffell, I., Hawkes, A. D., Li, F., Grünewald, P., McDowall, W., and Ekins, P. (2015). Hydrogen and fuel cell technologies for heating: A review. *International Journal of Hydrogen Energy*, 40(5):2065 – 2083. doi: <https://doi.org/10.1016/j.ijhydene.2014.11.059>.
- [96] Staffell, I., Scamman, D., Abad, A. V., Balcombe, P., Dodds, P. E., Ekins, P., Shah, N., and Ward, K. R. (2019). The role of hydrogen and fuel cells in the global energy system. *Energy & Environmental Science*, 12(2):463–491.
- [97] Speirs, J., Balcombe, P., Johnson, E., Martin, J., Brandon, N., and Hawkes, A. (2018). A greener gas grid: What are the options. *Energy Policy*, 118:291 – 297. doi: <https://doi.org/10.1016/j.enpol.2018.03.069>.
- [98] Harding, R. (2019). Japan launches first liquid hydrogen carrier ship. URL <https://www.ft.com/content/8ae16d5e-1bd4-11ea-97df-cc63de1d73f4>. [Online; accessed 11-January-2020].
- [99] Jensen, S. H., Graves, C., Mogensen, M., Wendel, C., Braun, R., Hughes, G., Gao, Z., and Barnett, S. A. (2015). Large-scale electricity storage utilizing reversible solid oxide cells combined with underground storage of co₂ and ch₄. *Energy Environ. Sci.*, 8:2471–2479. doi: 10.1039/C5EE01485A.
- [100] Butera, G., Jensen, S. H., and Clausen, L. R. (2019). A novel system for large-scale storage of electricity as synthetic natural gas using reversible pressurized solid oxide cells. *Energy*, 166:738 – 754. doi: <https://doi.org/10.1016/j.energy.2018.10.079>.
- [101] Cannon, D., Brayshaw, D., Methven, J., Coker, P., and Lenaghan, D. (2015). Using reanalysis data to quantify extreme wind power generation statistics: A 33 year case study in great britain. *Renewable Energy*, 75:767 – 778. doi: <https://doi.org/10.1016/j.renene.2014.10.024>.
- [102] IEA (2020). World Energy Outlook 2020. Technical report, IEA.
- [103] Heide, D., Greiner, M., von Bremen, L., and Hoffmann, C. (2011). Reduced storage and balancing needs in a fully renewable european power system with excess wind and solar power generation. *Renewable Energy*, 36(9):2515 – 2523. doi: <https://doi.org/10.1016/j.renene.2011.02.009>.
- [104] Rasmussen, M. G., Andresen, G. B., and Greiner, M. (2012). Storage and balancing synergies in a fully or highly renewable pan-european power system. *Energy Policy*, 51:642 – 651. doi: <https://doi.org/10.1016/j.enpol.2012.09.009>.

- [105] Pleßmann, G., Erdmann, M., Hlusiak, M., and Breyer, C. (2014). Global energy storage demand for a 100% renewable electricity supply. *Energy Procedia*, 46:22 – 31. doi: <https://doi.org/10.1016/j.egypro.2014.01.154>.
- [106] Luderer, G., Leimbach, M., Bauer, N., Kriegler, E., Baumstark, L., Bertram, C., Giannousakis, A., Hilaire, J., Klein, D., Levesque, A., Mouratiadou, I., Pehl, M., Pietzcker, R., Piontek, F., Roming, N., Schultes, A., Schwanitz, V. J., and Strefler, J. (2020). Model Documentation - REMIND. Technical report, Potsdam Institute for Climate Impact Research, Potsdam, Germany. URL <https://www.iamcdocumentation.eu>.
- [107] Clack, C. T. M., Choukulkar, A., Coté, B., and Mckee, S. A. (2020). Why Local Solar For All Costs Less: A New Roadmap for the Lowest Cost Grid. Technical report, Vibrant Clean Energy, LLC.
- [108] Steinke, F., Wolfrum, P., and Hoffmann, C. (2013). Grid vs. storage in a 100% renewable europe. *Renewable Energy*, 50:826–832. doi: <https://doi.org/10.1016/j.renene.2012.07.044>.
- [109] IEA (2020). World Energy Investment 2020. Technical report, International Energy Agency.
- [110] IEA (2022). World Energy Investment 2022. Technical report, International Energy Agency.
- [111] U.S. DOE (2012). Large power transformers and the u.s. electric grid. Technical report, Office of Electricity Delivery and Energy Reliability, U.S. Department of Energy.
- [112] Vaillancourt, K. (2014). Electricity transmission and distribution. Technology brief e12, IEA ETSAP.
- [113] U.S. EIA (2018). Assessing hvdc transmission for impacts of non-dispatchable generation. Technical report, U.S. Energy Information Administration, U.S. Department of Energy.
- [114] Burkes, K., Cordaro, J., Keister, T., and Cheung, K. (2017). Solid state power substation roadmap. Technical report, Office of Electricity Delivery and Energy Reliability, U.S. Department of Energy.
- [115] Hall, K. L. (2012). Out of sight, out of mind 2012 - an updated study on the undergrounding of overhead power lines. Technical report, Edison Electric Institute.
- [116] MISO (2019). Transmission cost estimation guide. Technical report, The Midcontinent Independent System Operator, Inc.
- [117] Short, T. A. (2005). *Electric Power Distribution Equipment and Systems*. CRC Press.
- [118] WECC (2012). Capital costs for transmission and substations - recommendations for wecc transmission expansion planning. Technical report, Western Electric Coordinating Council.
- [119] Reed, L., Morgan, M. G., Vaishnav, P., and Erian Armanios, D. (2019). Converting existing transmission corridors to hvdc is an overlooked option for increasing transmission capacity. *Proceedings of the National Academy of Sciences*, 116(28):13879–13884. doi: 10.1073/pnas.1905656116.
- [120] Kramer, G. J. and Haigh, M. (2009). No quick switch to low-carbon energy. *Nature*, 462(7273):568.
- [121] IPCC (2022). Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.
- [122] Byers, E., Krey, V., Kriegler, E., Riahi, K., Schaeffer, R., Kikstra, J., Lamboll, R., Nicholls, Z., Sandstad, M., Smith, C., van der Wijst, K., Lecocq, F., Portugal-Pereira, J., Saheb, Y., Stromann, A., Winkler, H., Auer, C., Brutschin, E., Lepault, C., Müller-Casseres, E., Gidden, M., Huppmann, D., Kolp, P., Marangoni, G., Werning, M., Calvin, K., Guivarch, C., Hasegawa, T., Peters, G., Steinberger, J., Tavoni, M., van Vuuren, D., Al Khourdajie, A., Forster, P., Lewis, J., Meinshausen, M., Rogelj, J., Samset, B., and Skeie, R. (2022). Ar6 scenarios database. URL <https://doi.org/10.5281/zenodo.5886912>.
- [123] Thompson, P. (2012). The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing. *The Journal of Economic Perspectives*, 26(3):203–224.
- [124] Arrow, K. J. (1962). The economic implications of learning by doing. *The review of economic studies*, pages 155–173.

- [125] Mazzola, J. B. and McCardle, K. F. (1997). The stochastic learning curve: Optimal production in the presence of learning-curve uncertainty. *Operations Research*, 45(3):440–450.
- [126] Nordhaus, W. D. (2014). The perils of the learning model for modeling endogenous technological change. *The Energy Journal*, Volume 35(1). doi: 10.5547/01956574.35.1.1.
- [127] Way, R., Lafond, F., Lillo, F., Panchenko, V., and Farmer, J. D. (2019). Wright meets markowitz: How standard portfolio theory changes when assets are technologies following experience curves. *Journal of Economic Dynamics and Control*, 101:211–238. doi: <https://doi.org/10.1016/j.jedc.2018.10.006>.
- [128] Neij, L. (1997). Use of experience curves to analyse the prospects for diffusion and adoption of renewable energy technology. *Energy Policy*, 25(13):1099 – 1107. doi: [https://doi.org/10.1016/S0301-4215\(97\)00135-3](https://doi.org/10.1016/S0301-4215(97)00135-3).
- [129] Rubin, E. S., Azevedo, I. M., Jaramillo, P., and Yeh, S. (2015). A review of learning rates for electricity supply technologies. *Energy Policy*, 86:198 – 218. doi: <https://doi.org/10.1016/j.enpol.2015.06.011>.
- [130] Samadi, S. (2018). The experience curve theory and its application in the field of electricity generation technologies – a literature review. *Renewable and Sustainable Energy Reviews*, 82:2346 – 2364. doi: <https://doi.org/10.1016/j.rser.2017.08.077>.
- [131] Wene, C.-O. (2016). Future energy system development depends on past learning opportunities. *WIREs Energy and Environment*, 5(1):16–32. doi: 10.1002/wene.172.
- [132] Ferioli, F., Schoots, K., and van der Zwaan, B. (2009). Use and limitations of learning curves for energy technology policy: A component-learning hypothesis. *Energy Policy*, 37(7):2525 – 2535. doi: <https://doi.org/10.1016/j.enpol.2008.10.043>.
- [133] Wilson, C., Grubler, A., Bento, N., Healey, S., De Stercke, S., and Zimm, C. (2020). Granular technologies to accelerate decarbonization. *Science*, 368(6486):36–39. doi: 10.1126/science.aaz8060.
- [134] Pindyck, R. S. (1999). The long-run evolution of energy prices. *The Energy Journal*, 20(2):1–27.
- [135] Shafiee, S. and Topal, E. (2010). A long-term view of worldwide fossil fuel prices. *Applied Energy*, 87(3):988 – 1000. doi: <https://doi.org/10.1016/j.apenergy.2009.09.012>.
- [136] Farmer, J. D. and Lafond, F. (2016). How predictable is technological progress? *Research Policy*, 45(3): 647–665.
- [137] BP (2021). BP Statistical Review of World Energy 2021. Technical report, BP plc.
- [138] Lazard (2021). Lazard’s Levelized Cost of Energy Analysis Versions 2-15. Technical report, Lazard.
- [139] IRENA (2021). Renewable Power Generation Costs in 2020. Technical report, International Renewable Energy Agency.
- [140] McNerney, J., Farmer, J. D., and Trancik, J. E. (2011). Historical costs of coal-fired electricity and implications for the future. *Energy Policy*, 39(6):3042 – 3054. doi: <https://doi.org/10.1016/j.enpol.2011.01.037>.
- [141] Colpier, U. C. and Cornland, D. (2002). The economics of the combined cycle gas turbine—an experience curve analysis. *Energy Policy*, 30(4):309 – 316. doi: [https://doi.org/10.1016/S0301-4215\(01\)00097-0](https://doi.org/10.1016/S0301-4215(01)00097-0).
- [142] Koomey, J. and Hultman, N. E. (2007). A reactor-level analysis of busbar costs for us nuclear plants, 1970–2005. *Energy Policy*, 35(11):5630 – 5642. doi: <https://doi.org/10.1016/j.enpol.2007.06.005>.
- [143] Nemet, G. F. (2006). Beyond the learning curve: factors influencing cost reductions in photovoltaics. *Energy Policy*, 34(17):3218 – 3232. doi: <https://doi.org/10.1016/j.enpol.2005.06.020>.
- [144] Glenk, G., Meier, R., and Reichelstein, S. (2021). Cost dynamics of clean energy technologies. *Schmalenbach Journal of Business Research*, 73(2):179–206. doi: 10.1007/s41471-021-00114-8.
- [145] Smil, V. (2016). *Energy Transitions, Global and National Perspectives*, 2nd Edition. ABC-CLIO.

- [146] Wisser, R., Jenni, K., Seel, J., Baker, E., Hand, M., Lantz, E., and Smith, A. (2016). Expert elicitation survey on future wind energy costs. *Nature Energy*, 1(10):16135. doi: 10.1038/nenergy.2016.135.
- [147] FRED (2022). Economic research at the federal reserve bank of st. louis. URL <https://fred.stlouisfed.org/>. [Online; accessed 23-April-2022].
- [148] Vaillancourt, K. (2014). Coal mining and logistics. Technology brief p07/08, IEA ETSAP.
- [149] Fouquet, R. (2008). *Heat, Power and Light: revolutions in energy services*. Edward Elgar Publications.
- [150] Seebregts, A. J. (2010). Gas-fired power. Technology brief e02, IEA ETSAP.
- [151] Lovering, J. R., Yip, A., and Nordhaus, T. (2016). Historical construction costs of global nuclear power reactors. *Energy Policy*, 91:371 – 382. doi: <https://doi.org/10.1016/j.enpol.2016.01.011>.
- [152] Koomey, J., Hultman, N. E., and Grubler, A. (2017). A reply to “historical construction costs of global nuclear power reactors”. *Energy Policy*, 102:640 – 643. doi: <https://doi.org/10.1016/j.enpol.2016.03.052>.
- [153] Gilbert, A., Sovacool, B. K., Johnstone, P., and Stirling, A. (2017). Cost overruns and financial risk in the construction of nuclear power reactors: A critical appraisal. *Energy Policy*, 102:644 – 649. doi: <https://doi.org/10.1016/j.enpol.2016.04.001>.
- [154] Sovacool, B. K., Gilbert, A., and Nugent, D. (2014). Risk, innovation, electricity infrastructure and construction cost overruns: Testing six hypotheses. *Energy*, 74:906 – 917. doi: <https://doi.org/10.1016/j.energy.2014.07.070>.
- [155] Grubler, A. (2010). The costs of the french nuclear scale-up: A case of negative learning by doing. *Energy Policy*, 38(9):5174 – 5188. doi: <https://doi.org/10.1016/j.enpol.2010.05.003>.
- [156] Matsuo, Y. and Nei, H. (2019). An analysis of the historical trends in nuclear power plant construction costs: The japanese experience. *Energy Policy*, 124:180 – 198. doi: <https://doi.org/10.1016/j.enpol.2018.08.067>.
- [157] Boccard, N. (2014). The cost of nuclear electricity: France after fukushima. *Energy Policy*, 66:450 – 461. doi: <https://doi.org/10.1016/j.enpol.2013.11.037>.
- [158] Court of Audit (2012). The costs of the nuclear power sector. Thematic public report, French Court of Audit.
- [159] IRENA (2019). Renewable Power Generation Costs in 2018. Technical report, International Renewable Energy Agency.
- [160] Ansar, A., Flyvbjerg, B., Budzier, A., and Lunn, D. (2014). Should we build more large dams? the actual costs of hydropower megaproject development. *Energy Policy*, 69:43 – 56. doi: <https://doi.org/10.1016/j.enpol.2013.10.069>.
- [161] IEA (2017). *Energy Technology Perspectives 2017 - Catalysing Energy Technology Transformations*. Technical report, International Energy Agency.
- [162] Junginger, M., de Visser, E., Hjort-Gregersen, K., Koornneef, J., Raven, R., Faaij, A., and Turkenburg, W. (2006). Technological learning in bioenergy systems. *Energy Policy*, 34(18):4024 – 4041. doi: <https://doi.org/10.1016/j.enpol.2005.09.012>.
- [163] Schilling, M. A. and Esmundo, M. (2009). Technology s-curves in renewable energy alternatives: Analysis and implications for industry and government. *Energy Policy*, 37(5):1767 – 1781. doi: <https://doi.org/10.1016/j.enpol.2009.01.004>.
- [164] Messner, S. (1997). Endogenized technological learning in an energy systems model. *Journal of Evolutionary Economics*, 7(3):291–313. doi: 10.1007/s001910050045.
- [165] Barreto, L. and Kypreos, S. (1999). Technological learning in energy models: Experience and scenario analysis with markal and the eris model prototype. Working Paper PSI Bericht Nr. 99-08, Paul Scherrer Institute.

- [166] Kouvaritakis, N., Soria, A., and Isoard, S. (2000). Modelling energy technology dynamics: methodology for adaptive expectations models with learning by doing and learning by searching. *International Journal of Global Energy Issues*, 14(1/2/3/4):104–115.
- [167] de Feber, M., Seebregts, A., and Smekens, K. (2002). Learning in clusters: Methodological issues and lock-out effects. Workshop Paper (International Energy Workshop), Energy Research Centre of the Netherlands.
- [168] Anderson, D. and Winne, S. (2004). Modelling innovation and threshold effects in climate change mitigation. Working Paper 59, Tyndall Cent. Clim. Change Res., Univ. East Anglia, Norwich, UK.
- [169] Baker, E., Fowle, M., Lemoine, D., and Reynolds, S. S. (2013). The Economics of Solar Electricity. *Annual Review of Resource Economics*, 5(1):387–426.
- [170] Kypreos, S. (2005). Modeling experience curves in merge (model for evaluating regional and global effects). *Energy*, 30(14):2721–2737. doi: <https://doi.org/10.1016/j.energy.2004.07.006>.
- [171] Bosetti, V., Carraro, C., Galeotti, M., Massetti, E., and Tavoni, M. (2006). A world induced technical change hybrid model. *The Energy Journal*, Hybrid Modeling(SI2):13–38. doi: 10.5547/ISSN0195-6574-EJ-VolSI2006-NoSI2-2.
- [172] Bosetti, V., Massetti, E., and Tavoni, M. (2007). The WITCH model. Structure, baseline, solutions. Working Paper 10.2007, Fondazione Eni Enrico Mattei.
- [173] Bosetti, V., Tavoni, M., De Cian, E., and Sgobbi, A. (2009). The 2008 WITCH Model: New Model Features and Baseline. Working Paper 85.2009, Fondazione Eni Enrico Mattei.
- [174] Rout, U. K., Blesl, M., Fahl, U., Remme, U., and Voß, A. (2009). Uncertainty in the learning rates of energy technologies: An experiment in a global multi-regional energy system model. *Energy Policy*, 37(11):4927–4942. doi: <https://doi.org/10.1016/j.enpol.2009.06.056>.
- [175] Luderer, G., Leimbach, M., Bauer, N., and Kriegler, E. (2011). Description of the ReMIND-R model. Technical report, Postdam Institute for Climate Impact Research, Potsdam, Germany.
- [176] Bibas, R. and Méjean, A. (2012). Potential and limitations of bioenergy for low carbon transitions. Working Paper No 42-2012, Centre International de Recherches sur l’Environnement et le Développement.
- [177] Bibas, R. and Méjean, A. (2014). Potential and limitations of bioenergy for low carbon transitions. *Climatic Change*, 123(3):731–761. doi: 10.1007/s10584-013-0962-6.
- [178] Criqui, P., Mima, S., Menanteau, P., and Kitous, A. (2015). Mitigation strategies and energy technology learning: An assessment with the poles model. *Technological Forecasting and Social Change*, 90:119–136. doi: <https://doi.org/10.1016/j.techfore.2014.05.005>.
- [179] Luderer, G., Leimbach, M., Bauer, N., Kriegler, E., Baumstark, L., Bertram, C., Giannousakis, A., Hilaire, J., Klein, D., Levesque, A., Mouratiadou, I., Pehl, M., Pietzcker, R., Piontek, F., Roming, N., Schultes, A., Schwanitz, V. J., and Strefler, J. (2015). Description of the REMIND model (version 1.6). Technical report, Postdam Institute for Climate Impact Research, Potsdam, Germany.
- [180] Emmerling, J., Drouet, L., Lara, A. R., Bevione, M., Berger, L., Bosetti, V., Carrara, S., De Cian, E., De Maere D’Aertrycke, G., Longden, T., Malpede, M., Marangoni, G., Sferra, F., Tavoni, M., Witajewski-Baltvilks, J., and Havlik, P. (2016). The WITCH 2016 Model — Documentation and Implementation of the Shared Socioeconomic Pathways. Working Paper 42.2016, Fondazione Eni Enrico Mattei.
- [181] Muratori, M., Ledna, C., McJeon, H., Kyle, P., Patel, P., Kim, S. H., Wise, M., Kheshgi, H. S., Clarke, L. E., and Edmonds, J. (2017). Cost of power or power of cost: A u.s. modeling perspective. *Renewable and Sustainable Energy Reviews*, 77:861–874. doi: <https://doi.org/10.1016/j.rser.2017.04.055>.
- [182] Creutzig, F., Agoston, P., Goldschmidt, J. C., Luderer, G., Nemet, G., and Pietzcker, R. C. (2017). The underestimated potential of solar energy to mitigate climate change. *Nature Energy*, 2(9):17140. doi: 10.1038/nenergy.2017.140.

- [183] Carrara, S., Bevione, M., de Boer, H. S., Gernaat, D., Mima, S., Pietzcker, R. C., and Tavoni, M. (2018). Exploring pathways of solar pv learning-by-doing in integrated assessment models. Post-print, HAL. URL <https://EconPapers.repec.org/RePEc:hal:journl:hal-01960855>.
- [184] IRENA (2021). Renewable Power Generation Costs in 2020. Technical report, International Renewable Energy Agency.
- [185] Nagy, B., Farmer, J. D., Bui, Q. M., and Trancik, J. E. (2013). Statistical basis for predicting technological progress. *PloS one*, 8(2):e52669.
- [186] Philipps, S. and Warmuth, W. (2022). Photovoltaics report. Technical report, Fraunhofer Institute for Solar Energy Systems, ISE, and PSE Projects GmbH.
- [187] IEA (2018). Global EV Outlook 2018 - Towards cross-modal electrification. Technical report, International Energy Agency.
- [188] Schmidt, O., Hawkes, A., Gambhir, A., and Staffell, I. (2017). The future cost of electrical energy storage based on experience rates. *Nature Energy*, 2(8):17110. doi: 10.1038/nenergy.2017.110.
- [189] Citibank (2016). Investment Themes in 2016.
- [190] Pillot, C. (2018). Current Status and Future Trends of the Global Li-ion Battery Market.
- [191] Clark, R. (2019). Solbat – challenges and opportunities from an industrial (materials) perspective. Faraday Institution meeting, February 21, 2019.
- [192] Lazard (2019). Lazard’s Levelized Cost of Storage (“LCOS”) Analysis Versions 1-5. Technical report, Lazard.
- [193] Kittner, N., Schmidt, O., Staffell, I., and Kammen, D. M. Chapter 8 - grid-scale energy storage. In Junginger, M. and Louwen, A., editors, *Technological Learning in the Transition to a Low-Carbon Energy System*, pages 119 – 143. Academic Press, 2020. ISBN 978-0-12-818762-3. doi: <https://doi.org/10.1016/B978-0-12-818762-3.00008-X>. URL <http://www.sciencedirect.com/science/article/pii/B978012818762300008X>.
- [194] Nykvist, B. and Nilsson, M. (2015). Rapidly falling costs of battery packs for electric vehicles. *Nature Climate Change*, 5(4):329–332. doi: 10.1038/nclimate2564.
- [195] Kittner, N., Lill, F., and Kammen, D. M. (2017). Energy storage deployment and innovation for the clean energy transition. *Nature Energy*, 2(9):17125. doi: 10.1038/nenergy.2017.125.
- [196] Kittner, N., Tsiropoulos, I., Tarvydas, D., Schmidt, O., Staffell, I., and Kammen, D. M. Chapter 9 - electric vehicles. In Junginger, M. and Louwen, A., editors, *Technological Learning in the Transition to a Low-Carbon Energy System*, pages 145 – 163. Academic Press, 2020. ISBN 978-0-12-818762-3. doi: <https://doi.org/10.1016/B978-0-12-818762-3.00009-1>. URL <http://www.sciencedirect.com/science/article/pii/B9780128187623000091>.
- [197] Pillot, C. (2019). The Rechargeable Battery Market and Main Trends 2018-2030.
- [198] Buttler, A. and Spliethoff, H. (2018). Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: A review. *Renewable and Sustainable Energy Reviews*, 82:2440 – 2454. doi: <https://doi.org/10.1016/j.rser.2017.09.003>.
- [199] Carmo, M., Fritz, D. L., Mergel, J., and Stolten, D. (2013). A comprehensive review on pem water electrolysis. *International Journal of Hydrogen Energy*, 38(12):4901 – 4934. doi: <https://doi.org/10.1016/j.ijhydene.2013.01.151>.
- [200] IRENA (2018). Hydrogen from renewable power: Technology outlook for the energy transition. Technical report, International Renewable Energy Agency.
- [201] Schoots, K., Ferioli, F., Kramer, G., and van der Zwaan, B. (2008). Learning curves for hydrogen production technology: An assessment of observed cost reductions. *International Journal of Hydrogen Energy*, 33(11):2630 – 2645. doi: <https://doi.org/10.1016/j.ijhydene.2008.03.011>.

- [202] Glenk, G. and Reichelstein, S. (2019). Economics of converting renewable power to hydrogen. *Nature Energy*, 4(3):216–222. doi: 10.1038/s41560-019-0326-1.
- [203] IRENA (2020). *Green Hydrogen Cost Reduction*. ISBN 97892922602956.
- [204] BNEF (2020). *Green Hydrogen: Time to Scale Up*.
- [205] IEA (2019). *World Energy Investment 2019*. Technical report, International Energy Agency.
- [206] EIA (2019). *Annual coal report 2018*. Technical report, U.S. Energy Information Administration.
- [207] Riahi, K., van Vuuren, D. P., Kriegler, E., Edmonds, J., O’Neill, B. C., Fujimori, S., Bauer, N., Calvin, K., Dellink, R., Fricko, O., Lutz, W., Popp, A., Cuaresma, J. C., KC, S., Leimbach, M., Jiang, L., Kram, T., Rao, S., Emmerling, J., Ebi, K., Hasegawa, T., Havlik, P., HumpenÄ¶nder, F., Silva, L. A. D., Smith, S., Stehfest, E., Bosetti, V., Eom, J., Gernaat, D., Masui, T., Rogelj, J., Strefler, J., Drouet, L., Krey, V., Luderer, G., Harmsen, M., Takahashi, K., Baumstark, L., Doelman, J. C., Kainuma, M., Klimont, Z., Marangoni, G., Lotze-Campen, H., Obersteiner, M., Tabeau, A., and Tavoni, M. (2017). The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, 42:153–168. doi: 10.1016/j.gloenvcha.2016.05.009.
- [208] Rogelj, J., Popp, A., Calvin, K. V., Luderer, G., Emmerling, J., Gernaat, D., Fujimori, S., Strefler, J., Hasegawa, T., Marangoni, G., Krey, V., Kriegler, E., Riahi, K., van Vuuren, D. P., Doelman, J., Drouet, L., Edmonds, J., Fricko, O., Harmsen, M., Havlík, P., HumpenÄ¶nder, F., Stehfest, E., and Tavoni, M. (2018). Scenarios towards limiting global mean temperature increase below 1.5 °c. *Nature Climate Change*, 8(4):325–332. doi: 10.1038/s41558-018-0091-3.
- [209] Gidden, M. J., Riahi, K., Smith, S. J., Fujimori, S., Luderer, G., Kriegler, E., van Vuuren, D. P., van den Berg, M., Feng, L., Klein, D., Calvin, K., Doelman, J. C., Frank, S., Fricko, O., Harmsen, M., Hasegawa, T., Havlik, P., Hilaire, J., Hoesly, R., Horing, J., Popp, A., Stehfest, E., and Takahashi, K. (2019). Global emissions pathways under different socioeconomic scenarios for use in cmip6: a dataset of harmonized emissions trajectories through the end of the century. *Geoscientific Model Development*, 12(4):1443–1475. doi: 10.5194/gmd-12-1443-2019.
- [210] van Vuuren, D. P., Stehfest, E., Gernaat, D. E., Doelman, J. C., van den Berg, M., Harmsen, M., de Boer, H. S., Bouwman, L. F., Daioglou, V., Edelenbosch, O. Y., Girod, B., Kram, T., Lassaletta, L., Lucas, P. L., van Meijl, H., MÄ¼ller, C., van Ruijven, B. J., van der Sluis, S., and Tabeau, A. (2017). Energy, land-use and greenhouse gas emissions trajectories under a green growth paradigm. *Global Environmental Change*, 42:237–250. doi: 10.1016/j.gloenvcha.2016.05.008.
- [211] Fricko, O., Havlik, P., Rogelj, J., Klimont, Z., Gusti, M., Johnson, N., Kolp, P., Strubegger, M., Valin, H., Amann, M., Ermolieva, T., Forsell, N., Herrero, M., Heyes, C., Kindermann, G., Krey, V., McCol-lum, D. L., Obersteiner, M., Pachauri, S., Rao, S., Schmid, E., Schoepp, W., and Riahi, K. (2017). The marker quantification of the shared socioeconomic pathway 2: A middle-of-the-road scenario for the 21st century. *Global Environmental Change*, 42:251–267. doi: 10.1016/j.gloenvcha.2016.06.004.
- [212] Fujimori, S., Hasegawa, T., Masui, T., Takahashi, K., Herran, D. S., Dai, H., Hijioka, Y., and Kainuma, M. (2017). SSP3: AIM implementation of shared socioeconomic pathways. *Global Environmental Change*, 42:268–283. doi: 10.1016/j.gloenvcha.2016.06.009.
- [213] Kriegler, E., Bauer, N., Popp, A., HumpenÄ¶nder, F., Leimbach, M., Strefler, J., Baumstark, L., Bodirsky, B. L., Hilaire, J., Klein, D., Mouratiadou, I., Weindl, I., Bertram, C., Dietrich, J.-P., Luderer, G., Pehl, M., Pietzcker, R., Piontek, F., Lotze-Campen, H., Biewald, A., Bonsch, M., Giannousakis, A., Kreiden-weis, U., MÄ¼ller, C., Rolinski, S., Schultes, A., Schwanitz, J., Stevanovic, M., Calvin, K., Emmer-ling, J., Fujimori, S., and Edenhofer, O. (2017). Fossil-fueled development (SSP5): An energy and resource intensive scenario for the 21st century. *Global Environmental Change*, 42:297–315. doi: 10.1016/j.gloenvcha.2016.05.015.
- [214] Calvin, K., Bond-Lamberty, B., Clarke, L., Edmonds, J., Eom, J., Hartin, C., Kim, S., Kyle, P., Link, R., Moss, R., McJeon, H., Patel, P., Smith, S., Waldhoff, S., and Wise, M. (2017). The SSP4: A world of deep-ening inequality. *Global Environmental Change*, 42:284–296. doi: 10.1016/j.gloenvcha.2016.06.010.

- [215] Masanet, E., Shehabi, A., Lei, N., Smith, S., and Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481):984–986. doi: 10.1126/science.aba3758.
- [216] Alberth, S. (2008). Forecasting technology costs via the experience curve—myth or magic? *Technological Forecasting and Social Change*, 75(7):952–983.
- [217] Lafond, F., Greenwald, D., and Farmer, J. D. (2020). Can stimulating demand drive costs down? world war ii as a natural experiment. INET Working Paper.
- [218] Pozzi, A. and Schivardi, F. (2016). Demand or productivity: what determines firm growth? *The RAND Journal of Economics*, 47(3):608–630. doi: <https://doi.org/10.1111/1756-2171.12142>.
- [219] Krey, V. (2014). Global energy-climate scenarios and models: a review. *WIREs Energy and Environment*, 3(4):363–383. doi: <https://doi.org/10.1002/wene.98>.
- [220] Bhattacharyya, S. C. and Timilsina, G. R. (2010). A review of energy system models. *International Journal of Energy Sector Management*.
- [221] Messner, S. and Strubegger, M. (1995). User’s guide for message iii.
- [222] Loulou, R., Goldstein, G., Noble, K., et al. (2004). Documentation for the markal family of models. *Energy Technology Systems Analysis Programme*, pages 65–73.
- [223] Loulou, R., Remme, U., Kanudia, A., Lehtila, A., and Goldstein, G. (2005). Documentation for the times model part ii. *Energy Technology Systems Analysis Programme*.
- [224] Howells, M., Rogner, H., Strachan, N., Heaps, C., Huntington, H., Kypreos, S., Hughes, A., Silveira, S., DeCarolis, J., Bazillian, M., et al. (2011). Osemosys: the open source energy modeling system: an introduction to its ethos, structure and development. *Energy Policy*, 39(10):5850–5870.
- [225] DeCarolis, J., Daly, H., Dodds, P., Keppo, I., Li, F., McDowall, W., Pye, S., Strachan, N., Trutnevyte, E., Usher, W., et al. (2017). Formalizing best practice for energy system optimization modelling. *Applied energy*, 194:184–198.
- [226] Keramidas, K., Kitous, A., Després, J., and Schmitz, A. (2017). Poles-jrc model documentation. Publications Office of the European Union.
- [227] Heaps, C. (2008). An introduction to leap. Stockholm Environment Institute, pages 1–16.
- [228] Grant, N., Hawkes, A., Napp, T., and Gambhir, A. (2020). The appropriate use of reference scenarios in mitigation analysis. *Nature Climate Change*, 10(7):605–610. doi: 10.1038/s41558-020-0826-9.
- [229] Ellenbeck, S. and Lilliestam, J. (2019). How modelers construct energy costs: discursive elements in energy system and integrated assessment models. *Energy Research & Social Science*, 47:69–77.
- [230] Trutnevyte, E. (2016). Does cost optimization approximate the real-world energy transition? *Energy*, 106:182–193.
- [231] Berglund, C. and Söderholm, P. (2006). Modeling technical change in energy system analysis: analyzing the introduction of learning-by-doing in bottom-up energy models. *Energy Policy*, 34(12):1344 – 1356. doi: <https://doi.org/10.1016/j.enpol.2004.09.002>.
- [232] Anandarajah, G. and McDowall, W. Multi-cluster technology learning in times: a transport sector case study with tiam-ucl. In *Informing energy and climate policies using energy systems models*, pages 261–278. Springer, 2015.
- [233] Heuberger, C. F., Rubin, E. S., Staffell, I., Shah, N., and Mac Dowell, N. (2017). Power capacity expansion planning considering endogenous technology cost learning. *Applied Energy*, 204:831–845.
- [234] Loulou, R. and Labriet, M. (2008). Etsap-tiam: the times integrated assessment model part i: Model structure. *Computational Management Science*, 5(1-2):7–40.
- [235] Pietzcker, R. C., Stetter, D., Manger, S., and Luderer, G. (2014). Using the sun to decarbonize the power sector: The economic potential of photovoltaics and concentrating solar power. *Applied Energy*, 135: 704–720.

- [236] Köhler, J., Grubb, M., Popp, D., and Edenhofer, O. (2006). The transition to endogenous technical change in climate-economy models: A technical overview to the innovation modeling comparison project. *Energy Journal*, 27(SPEC. ISS. MAR.):17–55. doi: 10.5547/issn0195-6574-ej-volsi2006-nosi1-2.
- [237] Kahouli-Brahmi, S. (2008). Technological learning in energy-environment-economy modelling: A survey. *Energy Policy*, 36(1):138–162. doi: 10.1016/j.enpol.2007.09.001.
- [238] Sano, F., Akimoto, K., Homma, T., and Tomoda, T. (2006). Analysis of technological portfolios for CO2 stabilizations and effects of technological changes. *Energy Journal*, 27(SPEC. ISS. MAR.):141–161. doi: 10.5547/issn0195-6574-ej-volsi2006-nosi1-6.
- [239] Energiewende, A. (2015). Current and future cost of photovoltaics—long-term scenarios for market development system prices and lcoe of utility-scale pv systems. Report by Fraunhofer Institute for Solar Energy Systems on behalf of Agora Energiewende.
- [240] Vivero-Pol, J. L. (2017). Food as commons or commodity? Exploring the links between normative valuations and agency in food transition. *Sustainability (Switzerland)*, 9(3). doi: 10.3390/su9030442.
- [241] Lang, S. Comparing the Results of Renewable Energy Auctions A Model-Based Analysis of Recent Trends from Around the World. PhD thesis, University of Cambridge, 2018.
- [242] Ondraczek, J., Komendantova, N., and Patt, A. (2015). WACC the dog: The effect of financing costs on the levelized cost of solar PV power. *Renewable Energy*, 75:888–898. doi: 10.1016/j.renene.2014.10.053.
- [243] Dupont, E., Koppelaar, R., and Jeanmart, H. (2018). Global available wind energy with physical and energy return on investment constraints. *Applied Energy*, 209(July 2017):322–338. doi: 10.1016/j.apenergy.2017.09.085.
- [244] Dupont, E., Koppelaar, R., and Jeanmart, H. (2020). Global available solar energy under physical and energy return on investment constraints. *Applied Energy*, 257:113968. doi: 10.1016/j.apenergy.2019.113968.
- [245] World Bank (2020). Solar Photovoltaic Power Potential by Country. Technical report, World Bank. URL <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/466331592817725242/global-photovoltaic-power-potential-by-country>.
- [246] IRENA (2020). Renewable Power Generation Costs in 2019. Technical report, International Renewable Energy Agency.
- [247] Jadun, P., McMillan, C., Stenberg, D., Muratori, M., Vimmerstedt, L., and Mai, T. (2017). Electrification Futures Study : End-Use Electric Technology Cost and. Technical report, National Renewable Energy Lab.
- [248] Eurek, K., Sullivan, P., Gleason, M., Hettinger, D., Heimiller, D., and Lopez, A. (2017). An improved global wind resource estimate for integrated assessment models. *Energy Economics*, 64:552–567. doi: 10.1016/j.eneco.2016.11.015.
- [249] Geels, F. W. (2019). Socio-technical transitions to sustainability: a review of criticisms and elaborations of the Multi-Level Perspective. ISSN 18773435.
- [250] Geels, F. W., Sovacool, B. K., Schwanen, T., and Sorrell, S. (2017). Sociotechnical transitions for deep decarbonization. *Science*, 357(6357):1242–1244. doi: 10.1126/science.aao3760.
- [251] Geels, F. W. (2002). Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. *Research Policy*, 31(8-9):1257–1274.
- [252] Grubler, A., Wilson, C., and Nemet, G. (2016). Apples, oranges, and consistent comparisons of the temporal dynamics of energy transitions. *Energy Research and Social Science*, 22:18–25. doi: 10.1016/j.erss.2016.08.015.

- [253] Geels, F. W., Sovacool, B. K., Schwanen, T., and Sorrell, S. (2017). The Socio-Technical Dynamics of Low-Carbon Transitions. *Joule*, 1(3):463–479. doi: <https://doi.org/10.1016/j.joule.2017.09.018>.
- [254] Nelson, S. and Allwood, J. M. (2021). The technological and social timelines of climate mitigation: Lessons from 12 past transitions. *Energy Policy*, 152:112155. doi: <https://doi.org/10.1016/j.enpol.2021.112155>.
- [255] Hanley, E. S., Deane, J. P., and Gallachóir, B. P. (2018). The role of hydrogen in low carbon energy futures—A review of existing perspectives. *Renewable and Sustainable Energy Reviews*, 82(July 2017): 3027–3045. doi: 10.1016/j.rser.2017.10.034.
- [256] Hydrogen Council (2020). Path to hydrogen competitiveness: a cost perspective.
- [257] Hall, W., Spencer, T., Renjith, G., and Dayal, S. (2020). The Potential Role of Hydrogen in India. Technical report, The Energy Research Institute.
- [258] Hydrogen Council (2017). Hydrogen scaling up: A sustainable pathway for the global energy transition. *Hydrogen scaling up: A sustainable pathway for the global energy transition*.
- [259] Marcucci, A., Kypreos, S., and Panos, E. (2017). The road to achieving the long-term Paris targets: energy transition and the role of direct air capture. *Climatic Change*, 144(2):181–193. doi: 10.1007/s10584-017-2051-8.
- [260] Blanco, H. and Faaij, A. (2018). A review at the role of storage in energy systems with a focus on Power to Gas and long-term storage. ISSN 18790690.
- [261] IEA (2012). Energy Technology Perspectives 2012 Pathways to a Clean Energy System. *Energy Technology Perspectives 2012 Pathways to a Clean Energy System*, pages 1–12.
- [262] Ives, M., Righetti, L., Schiele, J., De Meyer, K., Hubble-Rose, L., Tieng, F., Kruitwagen, L., Tillmann-Morris, L., Wang, T., Way, R., and Hepburn, C. (2021). A new perspective on decarbonising the global energy system. Technical Report 21-04, Oxford University Smith School of Enterprise and the Environment. URL www.energychallenge.info.
- [263] Millar, J. R., Nicholls, Z. R., Friedlingstein, P., and Allen, M. R. (2017). A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions. *Atmospheric Chemistry and Physics*, 17(11):7213–7228. doi: 10.5194/acp-17-7213-2017.
- [264] Smith, C. J., Forster, P. M., Allen, M., Leach, N., Millar, R. J., Passerello, G. A., and Regayre, L. A. (2018). FAIR v1.3: A simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6):2273–2297. doi: 10.5194/gmd-11-2273-2018.
- [265] Mitchell, D., Heaviside, C., Vardoulakis, S., Huntingford, C., Masato, G., Guillod, B. P., Frumhoff, P., Bowery, A., Wallom, D., and Allen, M. (2016). Attributing human mortality during extreme heat waves to anthropogenic climate change. *Environmental Research Letters*, 11(7):074006.
- [266] Rahmstorf, S. and Coumou, D. (2011). Increase of extreme events in a warming world. *Proceedings of the National Academy of Sciences*, 108(44):17905–17909.
- [267] Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives*, 32(4):33–52.
- [268] GAO (2014). Regulatory impact analysis: Development of social cost of carbon estimates. Technical report, United States Government Accountability Office. URL <https://www.gao.gov/assets/gao-14-663.pdf>.
- [269] Rising, J. A., Taylor, C., Ives, M. C., and Ward, R. E. (2022). Challenges and innovations in the economic evaluation of the risks of climate change. *Ecological Economics*, 197:107437. doi: 10.1016/j.ecolecon.2022.107437.
- [270] Kikstra, J. S., Waidelich, P., Rising, J., Yumashev, D., Hope, C., and Brierley, C. M. (2021). The social cost of carbon dioxide under climate-economy feedbacks and temperature variability. *Environmental Research Letters*, 16(9):094037. doi: 10.1088/1748-9326/ac1d0b.

- [271] Weitzman, M. L. (2010). What is the "damages Function" for Global Warming-And What Difference Might It Make? *Climate Change Economics*, 1(1):57–69. doi: 10.1142/S2010007810000042.
- [272] Cai, Y., Lenton, T. M., and Lontzek, T. S. (2016). Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction. *Nature Climate Change* 2016 6:5, 6(5):520–525. doi: 10.1038/nclimate2964.
- [273] Ricke, K., Drouet, L., Caldeira, K., and Tavoni, M. (2018). Country-level social cost of carbon. *Nature Climate Change*, 8(10):895–900.
- [274] Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523.
- [275] Smith, A. (2013). *The Climate Bonus. Co-benefits of Climate Policy*. Taylor & Francis Group Ltd.
- [276] IHME (2020). *Global Burden of Disease Study 2019*. Technical report, Institute for Health Metrics and Evaluation (IHME).
- [277] Rentschler, J. E. (2013). Oil price volatility, economic growth and the hedging role of renewable energy. *World Bank Policy Research Working Paper* 6603.
- [278] Generation Foundation (2013). *Stranded Carbon Assets*. Technical report, Generation Foundation. URL <http://www.phys.uri.edu/nigh/FFRI/pdf-generation-foundation-stranded-carbon-assets-v1.pdf> <http://genfound.org/media/pdf-generation-foundation-stranded-carbon-assets-v1.pdf>.
- [279] Caldecott, B. (2017). Introduction to special issue: stranded assets and the environment. *Journal of Sustainable Finance and Investment*, 7(1):1–13. doi: 10.1080/20430795.2016.1266748.
- [280] Tong, D., Zhang, Q., Zheng, Y., Caldeira, K., Shearer, C., Hong, C., Qin, Y., and Davis, S. J. (2019). Committed emissions from existing energy infrastructure jeopardize 1.5 °C climate target. ISSN 14764687.
- [281] Pfeiffer, A., Hepburn, C., Vogt-Schilb, A., and Caldecott, B. (2018). Committed emissions from existing and planned power plants and asset stranding required to meet the Paris Agreement. *Environmental Research Letters*, 13(5). doi: 10.1088/1748-9326/aabc5f.
- [282] Global Energy Monitor (2022). *Global Coal Plant Tracker*. URL <https://globalenergymonitor.org/projects/global-coal-plant-tracker/>.
- [283] Schumpeter, J. A. (1942). *Capitalism, Socialism, and Democracy*.