Institute for
New Economic Thinking
AT THE OXFORD MARTIN SCHOOL

Estimating the costs of energy transition scenarios using probabilistic forecasting methods

Rupert Way, Penny Mealy and J. Doyne Farmer

26th December 2020

INET Oxford Working Paper No. 2021-01

Estimating the costs of energy transition scenarios using probabilistic forecasting methods

Rupert Way[a,b], Penny Mealy[a,b,c] and J. Doyne Farmer[a,d,e]

[a]*Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, UK*
[b]*Smith School of Enterprise and the Environment, University of Oxford, Oxford, UK*
[c]*Bennett Institute for Public Policy, University of Cambridge, UK*
[d]*Mathematical Institute, University of Oxford, Oxford, UK*
[e]*Santa Fe Institute, Santa Fe, New Mexico, USA*

December 26, 2020

Abstract

We evaluate the cost of four different scenarios for the global energy system from 2020 to 2070 using an empirically validated technology forecasting method based on an expansive historical dataset. A no-transition scenario that maintains the current energy mix provides a benchmark. Under a rapid transition scenario, solar photovoltaics and wind are quickly deployed using batteries for short-term storage. Hydrogen-based fuels are used for long-term storage and non-electrifiable applications. Energy prices become lower than historical averages after 2030 and considerably lower after 2050. This yields an expected net present *saving* at any sensible discount rate; at 4% for example, we predict savings of $5.6 trillion. In contrast, a slower transition is more expensive, while a nuclear scenario is substantially more expensive.

Addressing climate change urgently requires a major transformation of our global energy system in order to eliminate fossil fuel greenhouse gas emissions. An important barrier to progress has been the estimated cost. William Nordhaus, who was awarded the Nobel Prize in economics in 2018, estimated that the costs of transitioning are likely to be so high that the 'optimal' future trajectory of humanity would be to incur 3.5°C warming by the end of the century (Nordhaus 2017). Other studies based on more sophisticated energy-economy models are more optimistic, but the estimated costs are still substantial (IEA & IRENA 2017, IRENA 2018, McCollum et al. 2018, Grubler et al. 2018, Ram et al. 2019). These models often lack transparency, involve assumptions that are at odds with empirical evidence, and are rarely tested against historical data (DeCarolis et al. 2012, Farmer et al. 2015, Trutnevyte 2016, Krey et al. 2019).

Future technology costs play a large role in determining energy model results, so it is crucial to have reliable forecasts. Here we use empirically validated methods of technology forecasting based on historical cost declines in order to estimate the distribution of future costs of different global energy scenarios. The experience curve forecasting model we use contains nonlinearities that prevent us from applying standard optimization tools to find optimal scenarios, so instead we analyze four contrasting scenarios. The reduced-form energy system model we develop enables calibration with relevant, up-to-date data, and allows us to model the deployment of energy technologies by simply assuming they grow at their current rates for roughly the next decade and then taper off once they dominate the energy system. Our approach is probabilistic, placing bets on the technologies that are likely to respond with rapid future improvements.

While there is considerable uncertainty, costs of key renewable energy technologies are likely to continue dropping at their current rapid rates, and the expected net present cost of a rapid green energy transition based on renewables is likely to be substantially cheaper than continuing with a fossil fuel based system. Our assessment is based purely on the economic costs of the energy system and does not include analysis of future climate damages or other benefits that would be realized by transitioning to renewable energy sources, such as improved health outcomes from lower air pollution. As detailed further in the discussion, our estimates should consequently be seen as a conservative lower bound on the actual value of accelerating the green transition.

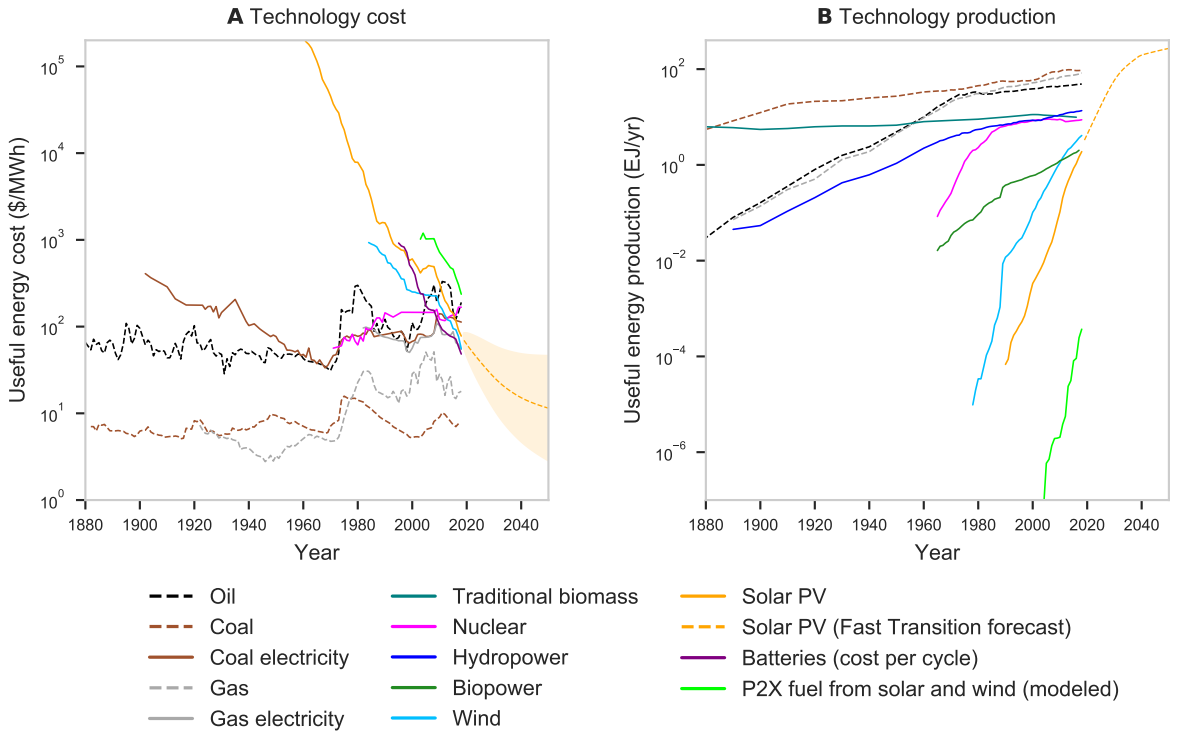

Figure 1: **Historical costs and production of key energy supply technologies**. A. *Inflation-adjusted useful energy costs as a function of time*. Fossil fuel prices have remained roughly constant for over a century, while solar photovoltaics, wind and batteries have experienced dramatic cost declines. The yellow band shows the 95% confidence level for a forecast of future solar PV prices in our fast green energy transition scenario. The median forecast cost is 11$/MWh in 2050, compared to current gas-fired electricity costs of 50$/MWh. Power-to-X (P2X) fuel costs assume historical electrolyzer costs and a 50-50 mix of solar and wind electricity. Note that *useful energy* accounts for the typical efficiency with which each technology provides energy services, so e.g. oil's particularly low efficiency accounts for its high useful energy cost. B. *Global useful energy consumption.* The provision of energy by solar photovoltaics has, on average, increased at 44% per year for the last 30 years, while wind has increased at 23% per year. These are just a few representative time series, for a full description of data and methods see the Supplementary Information (SI), Sections 1 and 6.

As illustrated in Figure 1, fossil fuel prices (oil, coal and natural gas) have been roughly constant for more than a century, making fluctuations of less than an order of magnitude, with no discernible inflation-adjusted trend. In contrast, the cost of solar photovoltaic energy has declined by more than three orders of magnitude since its first commercial application in 1958, dropping at about 10% per year for the last four decades. Wind and battery costs have also undergone substantial and steady exponential declines. Analysis of data of more than 50 diverse technologies demonstrates that technological cost trends are remarkably persistent (Farmer & Lafond 2016).

Here we construct a rapid green energy transition scenario based on technologies with favorable long-term trends, and show that with these choices, energy will likely soon become substantially cheaper than it has ever been, creating future savings that more than offset the expenditures required to implement the green energy transition.

Technology cost forecasting

In this paper we employ three simple models for the dynamics of technology costs. The first is a generalized form of *Wright's law*, which says that costs drop as a power law of cumulative production. (This relationship is also called an *experience curve* or *learning curve*, and cumulative production is also called *experience*). While Wright's law is not literally true, it is a very useful relationship that is widely applied to model endogenous learning in energy system models (Anandarajah et al. 2013, Heuberger et al. 2017, DeCarolis et al. 2017). Experience does not directly cause costs to drop, rather it is believed to be correlated with other factors that cause costs to drop, such as level of effort and R&D. Nonetheless, experience has the essential advantage of being relatively easy to measure, and Wright's law has been demonstrated to give useful out of sample forecasts (Alberth 2008, Nagy et al. 2013, Trancik et al. 2015, Lafond et al. 2018).

We use a first-difference stochastic form of Wright's law developed and tested by Lafond et al. (2018). 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 IID draw from a normal distribution. Then future costs are predicted using the iterative relationship

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

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 correlation parameter ρ characterizes the persistence of fluctuations in cost improvements.

A generalized form of *Moore's law* provides an alternative model in which technologies decrease in cost exponentially as a function of time. We use this in the form

$$\log c_t - \log c_{t-1} = -\mu + n_t + \theta n_{t-1}, \tag{2}$$

where $n_t \sim \mathcal{N}(0, \sigma_n^2)$ is an IID draw from a normal distribution. This relationship also has three parameters. For a given technology, the parameter μ is its exponential rate of cost improvement and the parameter σ_n^2 is the variability of this relationship. Again the correlation parameter θ characterizes the persistence of fluctuations in cost improvements.

If experience grows exponentially then Moore's law and Wright's law are equivalent, and the two produce forecasts of roughly equivalent quality (Nagy et al. 2013). However in general these two laws are not the same: If Moore's law is taken literally, then technological progress is inexorable – it will happen at the same rate regardless of policy. In contrast, if Wright's law is taken literally, then we can stimulate technological progress by stimulating production (or factors associated with it). This brings up the essential issue of causality: Do costs go down because experience increases, or does demand (and therefore cumulative production) increase because costs decrease?

Because we extrapolate existing trends, the issue of causality is not critical to our results. However, it is still worth noting that a recent study of technological change during WWII sheds some light on this question (Lafond et al. 2020). Military expenditures during WWII provide a good natural experiment because funding allocations were largely demand-driven rather than cost driven, so the direction of causality is clear. An exhaustive analysis of 675 different types of military equipment produced

by the US in WWII suggests that roughly 50% of cost decreases are associated with exogenous exponential trends (corresponding to Moore's law) and 50% due to experience (corresponding to Wright's law). We use Wright's law as our main specification, but in the SI we show that using Moore's law only strengthens our results, as in the rapid renewable scenario it causes prices to continue to drop rapidly even after solar and wind become dominant and deployment slows down.

If we attempt to apply the stochastic form of Wright's law given above to technologies that show relatively little historical cost improvement, such as fossil fuels, we find that $\omega \approx 0$, and Eq. (1) becomes a geometric random walk with no drift. This is a common model for financial time series, including tradable commodities such as oil or gas, and can be justified based on the efficient markets hypothesis. On short timescales (say ten years or less) this is a reasonable approximation, but over longer timescales it predicts too much volatility. Fossil fuel time series show mean reversion on longer timescales (Pindyck 1999), so are better captured by an AR(1) process of the form

$$\log c_t = \log c_{t-1} + \beta(\mu - \log c_{t-1}) + \epsilon_t, \qquad \text{with IID } \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \tag{3}$$

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 and β is the rate of mean reversion.

All three models have the important advantage that they provide estimates of the future distribution of costs, making it possible to assign a likelihood for cost to any given scenario. The distribution estimates using Moore's law and Wright's law have been carefully tested and shown to be accurate using historical data for 50 different technologies (Lafond et al. 2018). Similarly, a multivariate AR(1) model has previously been tested against historical fossil fuel prices and shown to perform well (Pindyck 1999). For simplicity, we use the univariate model (Eq. 3) and apply this for coal, oil, and gas. While coal-fired electricity and gas-fired electricity showed significant drops in cost for much of the twentieth century, in the long run their costs are increasingly dominated by fuel costs (McNerney et al. 2011), so we model these using the AR(1) specification as well. We use Wright's law for all other technologies. Although the historical cost trends for nuclear and hydroelectric power are flat or rising, we optimistically assume that the exponent $\omega \approx 0$ and σ_u is small so costs are close to constant.

Many previous studies assume floor costs (De Cian et al. 2016), i.e. a minimum level below which costs cannot fall. However, to the best of our knowledge no empirical evidence exists to support such an assumption. Indeed, in a forecasting study using empirical data for 50 technologies, there were no obvious signs in the data that floor costs may exist (Farmer & Lafond 2016). Given the lack of empirical evidence for imposing such restrictions, we do not include floor costs in our model.

Model design

Our model is designed to estimate the expected net present cost (NPC) of the transition from the current energy system to alternative energy systems, and to understand the most important factors influencing differences in cost. Our approach is based on two key design principles: 1) include only the minimal set of variables necessary to accurately represent the global energy system and its evolution under different scenarios, and 2) ensure all assumptions and dynamics are technically realistic and as

closely tied to empirical evidence as possible.

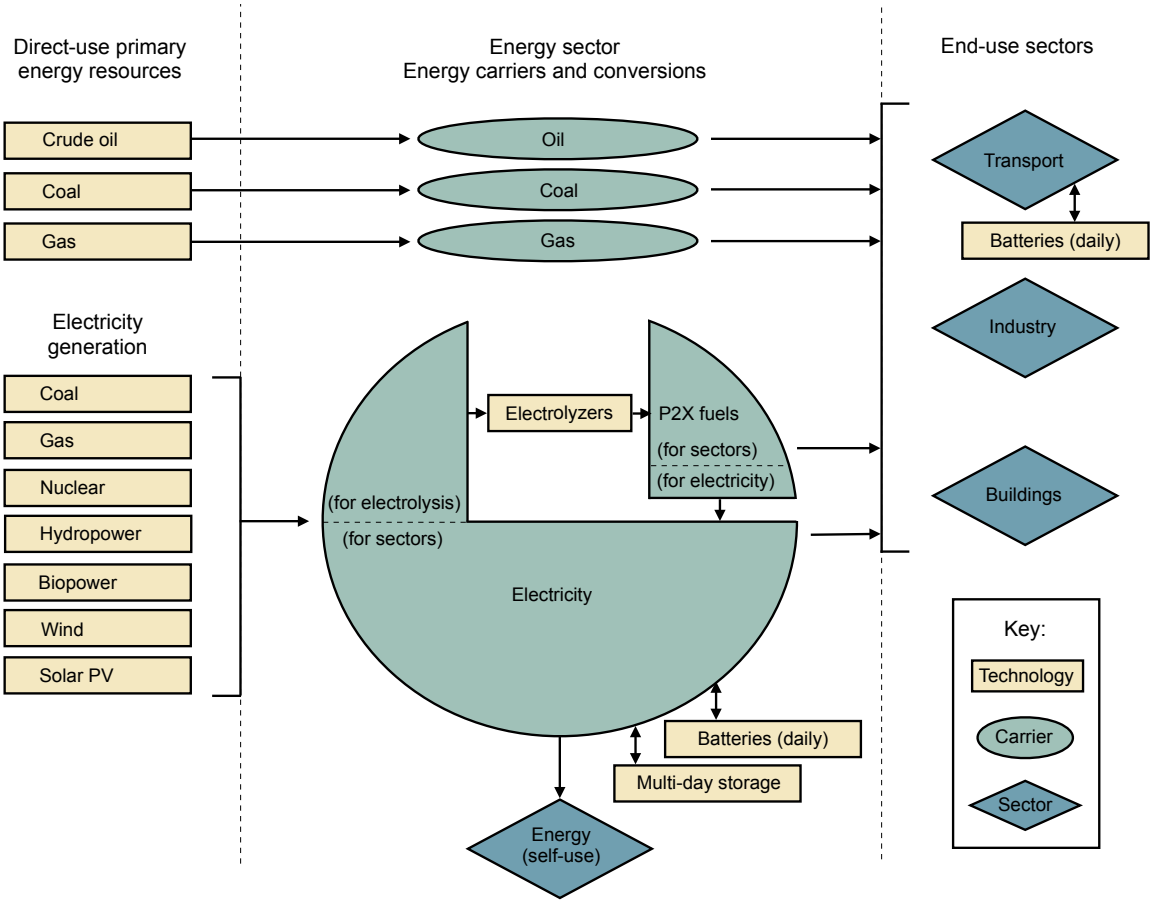

Figure 2: **Energy system model components**. We model 13 energy technologies (shown in yellow), which facilitate the supply, storage and conversion of energy, in the form of various energy carriers (shown in green). Energy sources are shown on the left and end-use sectors of the economy are shown on the right. Fossil fuels can be used directly for many purposes, but the renewable sources here only generate electricity, which can either be used directly, stored in batteries, or converted to P2X fuels and then used for purposes such as transport or industrial heat. All four scenarios supply the same amount of useful energy and satisfy the needs of each end-use sector.

Our model captures all major energy flows and key conversions from primary to final to useful energy. *Primary energy* is energy embodied in resources as they exist in nature. *Final energy* is the energy delivered to end-use sectors (transport, industry and buildings) and the energy sector itself. *Useful energy* is the portion of final energy used to perform energy services, such as heat, light and kinetic energy.

Figure 2 shows a stylized illustration of our model. Technologies (and their units) have been selected on the basis of both the empirical verifiability of their current and historical costs, which allows us to forecast their likely future improvement, and their ability to collectively approximate all investment and running costs of the energy system. We therefore exclude co-generation of heat, traditional biomass, petrochemical feedstock and energy sector fossil fuel consumption from our model, and cover 93% of the remaining energy system. Technologies that are currently minor players and have not exhibited significant historical cost improvements are also omitted. Thus we

exclude carbon capture and storage, tidal energy, solar thermal energy and geothermal energy. We also omit liquid biofuels because any significant contribution to the global energy system would have high environmental costs. (See the SI for a complete description of the model).

The efficiency with which sectors convert final energy to useful energy varies dramatically across sectors and energy carriers. We model the conversion process based on average final-to-useful energy conversion efficiency factors given in De Stercke (2014). In comparison to electricity, fossil fuels tend to have fairly large conversion losses, which means that significantly more final energy needs to be produced to obtain a given amount of useful energy. However, by switching to energy carriers with higher conversion efficiencies (e.g. moving to a fleet of electric vehicles in the transport sector), final energy consumption can be significantly reduced (Eyre 2019, Lovins et al. 2019). We make extensive use of efficiency savings in the development of our scenarios.

Scenarios

We construct four scenarios, subject to the constraint that they must all provide the same total useful energy and supply the needs of each end use sector (see Table 1). We take useful energy as a proxy for energy services and assume that total useful energy continues to grow at its current rate of 2% per year from now until 2070 when the scenarios end. Modeled scenarios include: (A) a no-transition scenario, in which the energy system remains similar to its current form and grows proportionally; (B) a fast transition scenario in which renewable energy technologies maintain their current high exponential growth rates for about a decade, replacing fossil fuels in about two decades, then slow to grow at 2% per year; long term energy storage is achieved with *Power-to-X* (P2X) fuels, i.e. by using electricity for hydrogen electrolysis and then making other fuels such as ammonia and methane as needed (Davis et al. 2018). (C) a slow renewable transition scenario, in which fossil fuels are phased out over a longer time horizon and (D) a nuclear scenario which is similar to the slow transition scenario, but with nuclear playing a dominant role in providing carbon-free energy.

A. *No Transition* The current energy mix prevails, meaning that the relative size of all energy sources are maintained approximately constant at their current values, with a pro rata increase of 2% per year. Final energy remains high because electrification does not occur.
B. *Fast Transition* Current high growth rates in clean energy renewable technologies prevail for a decade, then gradually relax back to the systemwide rate. Fossil fuels are quickly displaced from all sectors, and all essential fuel uses are replaced by P2X fuels. Solar and wind provide most of the energy, and reliable electricity is maintained by energy storage technologies based on batteries and P2X fuels.
C. *Slow Transition* A mix of the first two scenarios. Conversion to renewables is slow until 2050 then gathers pace, with solar and wind taking large shares. Far less P2X fuel is required because fossil fuels are not fully displaced from the system on this time scale.
D. *Slow Nuclear Transition* Nuclear takes the largest share of electricity generation. The pace of replacement is similar to C.

Table 1: Description of the four scenarios. In each case total useful energy grows at 2% per year and the energy carrier needs of all sectors are satisfied.

The endpoint of each scenario in 2070 is defined by the shares of technologies providing electricity generation and the shares of energy carriers providing final energy. All scenarios match in terms of their current shares in 2020. Transitions are implemented by assuming growth rates follow logistic (or "S") curves with a specified start and end point consistent with the growth of total useful energy. For the fast renewable transition, for example, we initialize the transition by matching current growth rates for solar, wind, batteries and P2X fuels. Details of all growth rates, timings and energy and carrier mixes for each scenario are given in the SI.

Are the postulated rates of increase of renewable energy capacity plausible? As seen in Figure 1, from 1990-2018, solar PV output grew at an average exponential rate of 44% per year and wind grew at a rate of 23% per year. We assume that this exponential growth continues for roughly a decade at roughly the same rate and then tapers off. We conservatively assume the initial growth rates are 32% for solar PV and 20% for wind. Although presently these technologies only produce 4% of useful energy, extrapolating these trends means that solar and wind take over 70% of all useful energy provision by 2050, leaving fossil fuels with a declining share of just 3%, with nuclear, hydropower and other sources making up the remainder. It is beyond the scope of this paper to carefully evaluate greenhouse gas reductions for each scenario. However, since the fast transition scenario quickly eliminates most fossil fuels, its associated emissions are well below those of the Paris agreement.

In order for renewables to become dominant it is essential to solve the intermittency problem. We assume intermittency issues are addressed using a mixture of batteries for short term storage and P2X fuels for long term storage. An overcapacity of solar and wind is built to perform hydrogen electrolysis when there is productive

weather. During spells of unproductive weather hydrogen is used to generate electricity using either fuel cells or turbines. Gas turbines are now being developed to burn 100% hydrogen and can be easily retrofit where appropriate (Staffell et al. 2019, Mitsubishi Hitachi Power Systems 2019), and in many cases hydrogen production can happen near power generation facilities.

Hydrogen electrolysis is the first stage for producing denser and more portable fuels, such as ammonia, that can be used to replace fossil fuels for end uses that cannot be accomplished with electricity or hydrogen. At current prices electricity is about 85% of the cost of making electrolytic hydrogen, and making hydrogen is 80% of the cost of making ammonia (IEA 2019, Brynolf et al. 2018), so these will both benefit substantially from decreasing electricity costs. We have added ample supply of batteries and electrolyzers so that renewables can provide all energy needs without extensive infrastructure modifications, e.g. without dramatic alterations to the grid.

Electricity for P2X fuels is a significant cost in the rapid transition scenario. Infrastructure costs per unit of P2X distribution and storage are currently 2-3 times that of fossil fuels; we pessimistically assume that unit costs remain constant even as the scale increases. However, because the needed deployment is 2-3 times less than that of existing fossil fuel infrastructure, we estimate the total infrastructure cost remains roughly constant. Investments in electricity smart grids and interconnection may provide cheaper solutions; if so this would only strengthen our results. (We do not use these because future costs are difficult to estimate).

Figure 3 shows how the energy system evolves in each scenario. Although total useful energy supplied is identical across scenarios, efficiency gains associated with electrification result in transition scenarios requiring less final energy than the no-transition scenario.

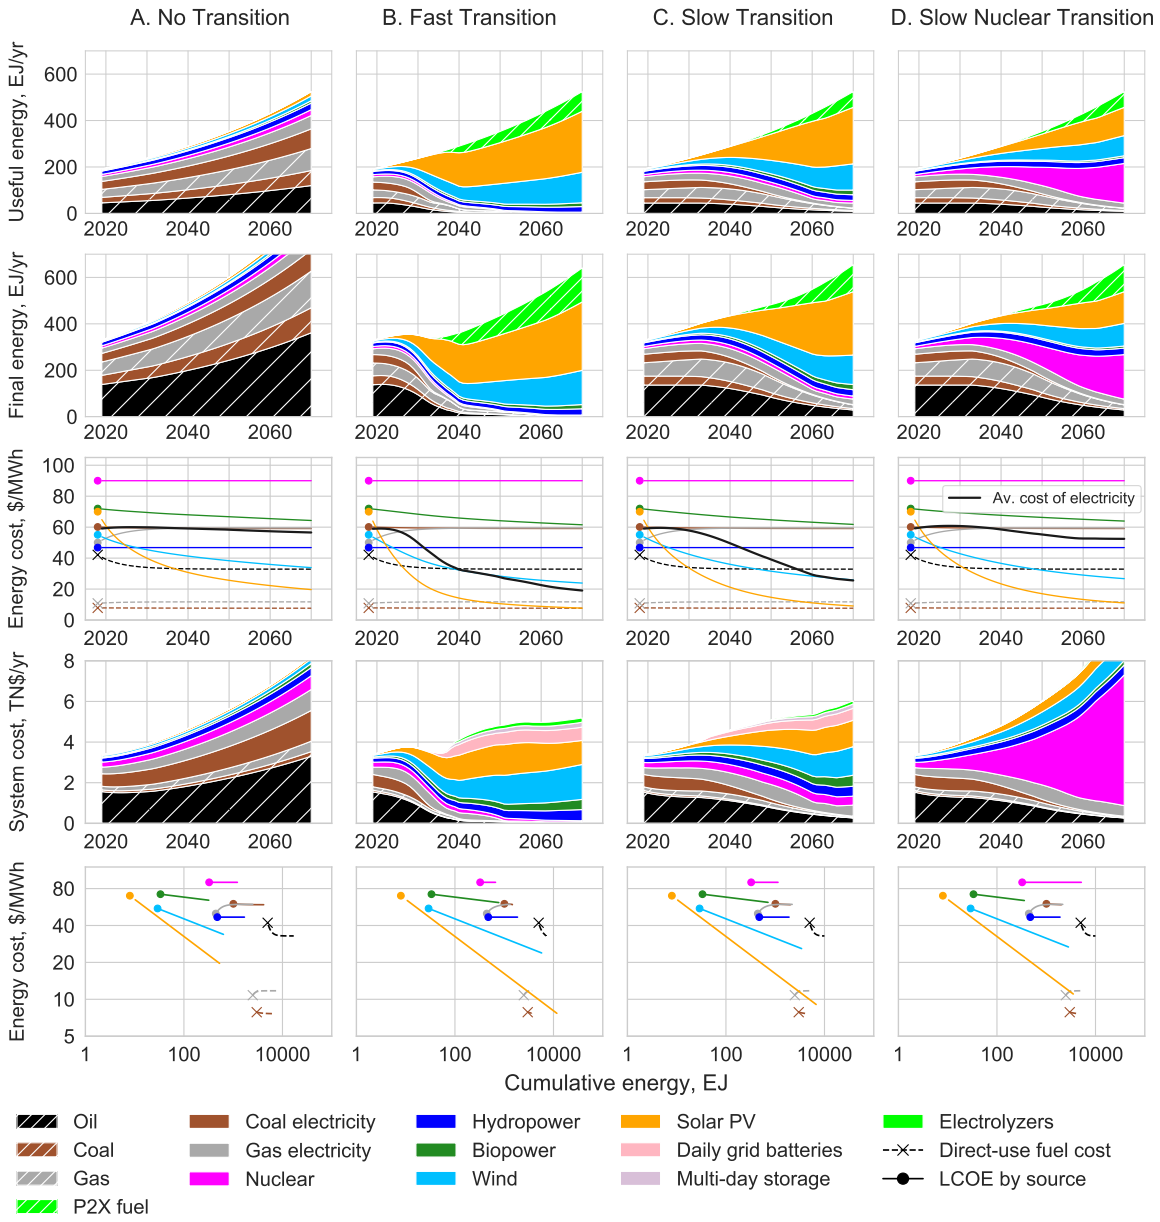

Figure 3: **Scenarios and their cost performance**. The four columns represent each of the scenarios. The rows are: (1) The useful energy provided by each technology and P2X fuel as a function of time. (2) The final energy provided by each technology and P2X fuel as a function of time. (3) The median unit energy cost of each technology as a function of time. (4) The median annual total spend on each technology as a function of time. (5) The median unit cost of each technology as a function of experience (i.e. cumulative production).

How much will each scenario cost?

We now apply the technology cost forecasting models described earlier to estimate the future costs of each scenario. Row 3 of Figure 3 shows how the median costs of energy technologies evolve over time in each scenario. Since fossil fuel costs are modeled with the AR(1) process, they quickly relax to fluctuate around their long-run historical averages and remain the same across all scenarios. Similarly, the model projects that costs for hydropower electricity and nuclear energy will remain fairly constant, in

accordance with their historical trends. Solar PV differs substantially across scenarios due to differences in its cumulative production and the corresponding costs under Wright's law.

New solar and wind installations are already the same cost or cheaper than new coal and gas installations. In the fast transition, the global median levelized cost of solar reaches the *operating costs* of fully-depreciated coal and gas-fired electricity (around 30$/MWh) in 2028, so in many locations it will be cheaper to install a new solar farm than to operate an existing gas power plant. By 2070 solar is estimated to cost roughly 8$/MWh, making it 77% cheaper than coal- and gas-fired electricity operating costs. In contrast, in the no-transition scenario, solar remains more expensive than coal- and gas-fired electricity operating costs until 2042.

For each scenario we aggregate cost estimates and exponentially discount future costs to calculate the expected net present cost of the global energy system. Figure 4A shows the relative expected net present cost for each scenario as a function of the discount rate, using the no-transition scenario as a benchmark. Previous analyses have suggested that whether it makes good economic sense to quickly transition to clean energy technologies depends on the discount rate. However, here we show a striking result: at all reasonable discount rates the fast renewable energy transition is likely to be substantially cheaper than the existing fossil-fuel based energy system (the cut off is at a discount rate of about 25%). Using the 1.4% discount rate advocated in the Stern Review (Stern 2007), for example, the expected net present saving is roughly $11 trillion. The median value, which gives a better indication of the net present saving likely to be realized in practice, is roughly $24 trillion. (The distribution of costs is roughly log-normal, so means and medians are substantially different).

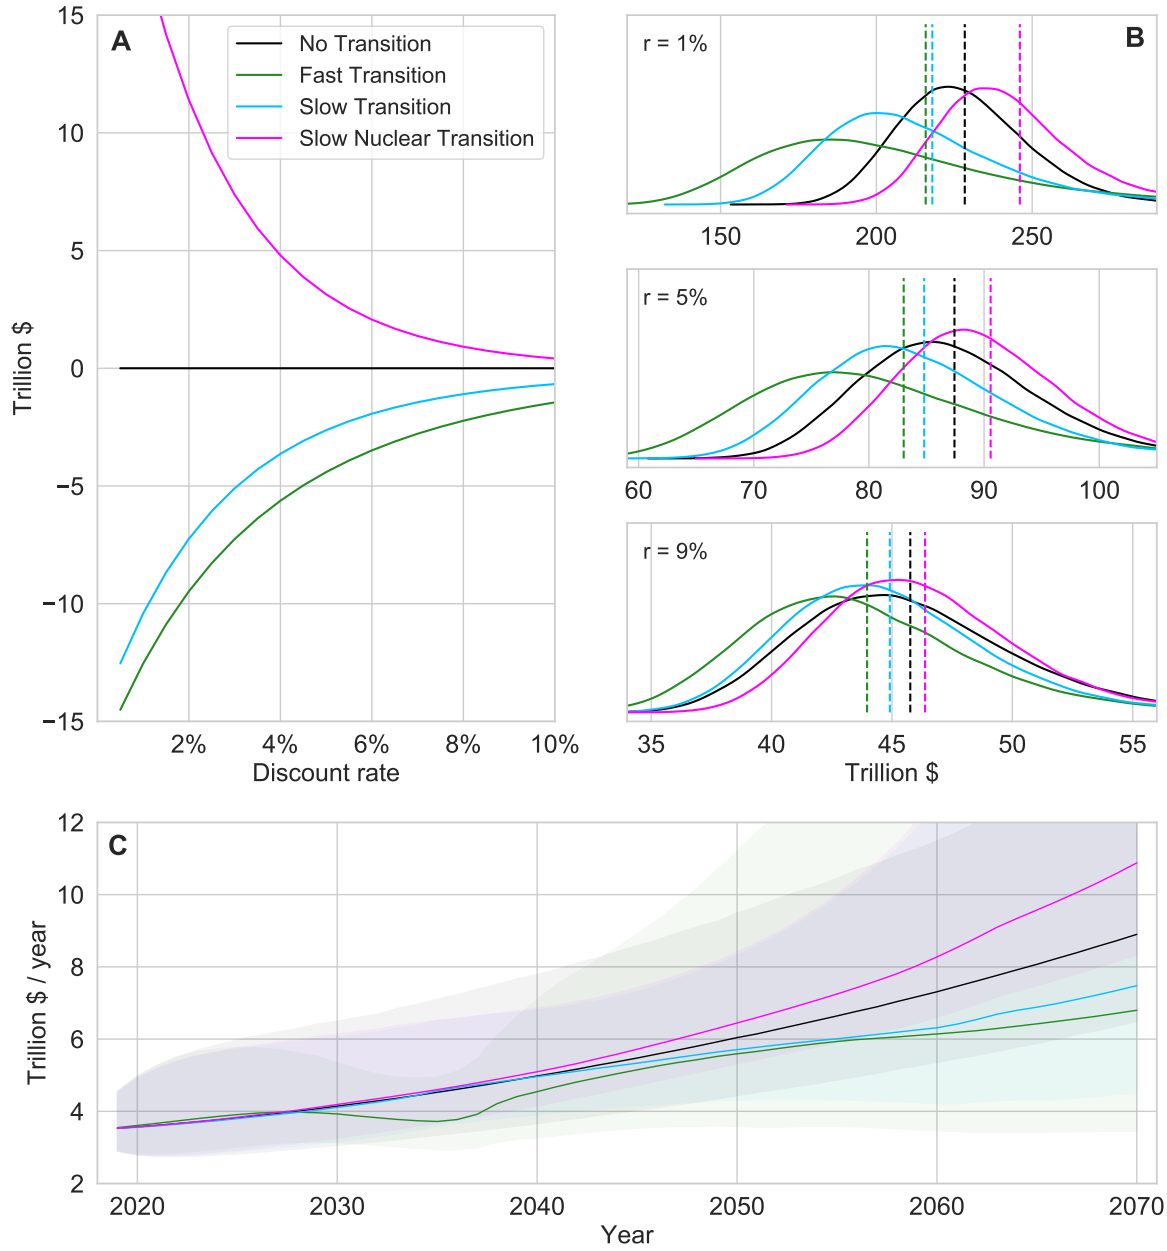

Figure 4: **Net present cost as a function of discount rate.** A. Expected (mean) net present costs for each scenario relative to the No Transition baseline. B. Distribution of present discounted costs for each scenario for three different discount rates, with dashed lines showing their expected values. C. Median annual system cost with shaded 95% confidence interval for each scenario.

The slow renewable transition also generates savings relative to no transition, though it is not as cheap as the fast transition. This is because the slow ramp up of renewables means that their savings are only realized much later, and the current high spending on fossil fuels continues for the next few decades. Unsurprisingly, given historical cost trends, the nuclear transition scenario is much more expensive than the others.

Figure 4B shows the projected distributions of present discounted costs for each scenario. The outcomes of the scenarios overlap due to the substantial uncertainty in future costs. These distributions can be used to estimate the likelihood that the rankings of any given pair of scenarios will change. As the discount rate decreases, due

to the increased uncertainty of the cost estimates in the far future, the distributions become broader. At the same time, due to the increase in the extent to which future benefits are felt, the difference in their expected net present cost also widens. These two effects counteract each other, making the results fairly insensitive to discount rate. For example, using a 2% discount rate gives a probability that the fast transition scenario will turn out to be cheaper than the no-transition scenario of 73%, while a 10% discount rate only increases this to 74%.

We have not systematically searched the space of possible scenarios, and we do not claim that the fast renewable scenario presented here is the cheapest possibility. Given the low cost of gas, it could be possible to achieve cheaper scenarios by using gas in place of P2X fuels in some applications, but this of course would not be a zero-emissions system. Similarly, while fossil fuel costs have not historically trended down, competition from renewables may force them down; this is feasible only at substantially reduced production levels where only the cheapest fossil fuel producers are competitive. This suggests that, while most of the fast renewable transition can be aligned with market forces, policies that discourage the use of fossil fuels will likely still be needed to fully decarbonize energy.

To enhance the credibility of our estimates we have been consistently conservative in estimating costs of renewables. Our systematic use of technologies with appropriate price histories means that in many cases we were forced to neglect solutions that many experts consider the most promising avenues of cost reduction, such as demand-side management of power grids and end-use efficiency (Lovins et al. 2019). We have considered global average energy costs, but in reality there is considerable variation by location, and we expect that future efforts will take advantage of this to achieve even cheaper solutions. As a result of these factors, it is likely that the future costs of the renewable transition will be substantially lower than those modeled here.

Discussion

Explanation of key assumptions

Is it reasonable to assume that the installed capacity of solar and wind can continue to grow at their current rates for the next decade, as they do in the fast renewable transition scenario? Based on the history of deployment of other energy technologies, Kramer & Haigh (2009) have 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, with no supporting analysis or validation. As we argue in the SI, the economics of solar and wind are much more favorable than that of previous technologies such as nuclear energy, and we see no compelling reason why these trends cannot continue for some time.

What about sunk costs? The existing energy system has sunk costs. Lifetimes of energy generation systems range from 25 to 50 years, meaning that on average about 2-4% of capacity needs replacing in any given year. When combined with the historical useful energy growth rate of 2% per year, this means that renewables can grow at a rate corresponding to roughly 5% of current energy demand without needing to retire existing capacity. This allows them to replace most of the existing energy system in 20 years and replace the remaining 5% within a few decades, as is done in the fast renewable transition scenario. In fact, if we are to meet climate objectives, significant

asset stranding is likely to occur in any case. For example, a recent study estimated that in order to stay within a 1.5 - 2°C global warming limit, around 20% of global power plant capacity is likely to become stranded (Pfeiffer et al. 2018). As mentioned earlier, after 2030, stranding gas and coal power plants by replacing them with solar PV and wind is likely to become cheaper than continuing to operate them in any case.

How do we know that increasing solar and wind capacity will continue to drive costs down? Our methods for estimating technological cost declines are based on extrapolating past trends, either as a function of time or experience. We do not provide a causal explanation. While it is possible that these trends will be broken, the analysis of 50 technologies by Farmer & Lafond (2016) indicates that this is rare.

What about storage? A substantial increase in energy storage is essential for nondispatchable energy sources to dominate the energy supply. As detailed in the SI, the situation for short term storage using batteries is favorable. From 1995 to 2018 the production of lithium ion batteries increased at 30% per year, while costs dropped at 12% per year, giving an experience exponent of 0.34, comparable to that of solar PV. Historically most of this production was in the electronics sector, but EVs and the power sector are likely to dominate future growth. Our fast transition scenario assumes growth continues at its current rate for the next decade then tapers off.

The most speculative part of our analysis relates to long term storage. Batteries can provide short term energy storage but they are not currently an effective solution for long term bulk storage and do not address the need for high energy density fuels. About 85% of the cost of electrolytic hydrogen is electricity, so it piggy-backs on the high progress rates seen in solar PV and wind: Lower future renewable costs automatically lead to lower P2X fuel costs. In addition, due to increased R&D efforts and gradual commercialisation, electrolyzers have seen substantial cost reductions over the last decade, and current short term forecasts from manufacturers indicate significant further declines in the next decade. Conversion of hydrogen to more manageable ammonia and methane will be necessary in many applications. These processes entail extra costs and energy losses, but they are already well understood and widely used in commercial applications, so present a viable pathway to decarbonisation.

Why is the forecast for nuclear so pessimistic? Couldn't we drive the costs of nuclear energy down? Nuclear power is a highly controversial topic that elicits emotionally charged views from both opponents and advocates. Costs vary considerably by country, but in almost every case, historical trends are not favorable. In many countries, such as France and the US, costs have risen substantially (Koomey & Hultman 2007, Grubler 2010). The most optimistic experience is for South Korea, where by some estimates costs have dropped by about 2% per year, but this is controversial (Koomey et al. 2017), and in any case much smaller than the 10% improvement rate of solar PV. Many advocates favor small modular reactors, but these have high initial capital costs, fight economies of scale based on fundamental physics, and have not yet been produced commercially, so are unlikely to be viable on the rapid time scale needed to combat climate change.

Why are the results so different from other models' results? Nonlinearities in the experience curve model prevent the use of standard optimization methods for finding optimal scenarios. However, large energy models generate optimal scenarios by applying such methods to hundreds or thousands of interacting technologies, so implementing experience curves without significant model redesign is infeasible. Instead they either use deterministic experience curves for a small number of technologies, or sim-

ply exogenous cost forecasts. In both cases technology growth is limited by internal modeling assumptions and historical growth rates are rarely a focus of the analysis, so whole regions of low-cost scenarios are left unexplored (Creutzig et al. 2017). The high geographical and technological resolution of large models also makes them hard to calibrate with current data. By using a much smaller model that places emphasis on accurately portraying historical technological trends — and allowing them to continue into the future — we obtain very different results, that we believe better reflect future prospects.

Additional benefits

In addition to the economic savings associated with rapidly driving energy costs down, the broader benefits of reducing green house gas emissions quickly and averting the worst impacts of climate change are manifest. The world is already experiencing more extreme weather events and adverse impacts that can be shown to be attributable to warming temperatures (Mitchell et al. 2016, Rahmstorf & Coumou 2011). Studies are also increasingly showing that the social cost carbon (SCC), which represents the present social value of damages from emitting CO_2, is significantly higher than was previously estimated (Auffhammer 2018). For example, incorporating the possibility of non-linear 'tipping points' in the climate system increases SCC estimates by a factor of three (i.e from $45 to $154 per tonne) (Cai et al. 2015), and a more recent study drawing on the latest climate projections and more robust macroeconomic modelling approaches finds that the global SCC stands at around $417 per ton (Ricke et al. 2018). As a point of comparison, given that we currently emit about 36 billion tons of carbon per year, under the no transition scenario at $100 per ton the SCC is about 3.6 trillion per year. With a 4% discount rate this corresponds to $90 trillion, i.e. it is comparable to the cost of the energy system itself (see Fig. 4).

There are also other important benefits of accelerating the green transition beyond those directly associated with reducing adverse climate change impacts. For example, the World Health Organization estimates that 4.2 million lives per year are lost to ambient air pollution, which is mostly caused by fossil fuels. 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 prices (Rentschler 2013). From a geopolitical point of view, the use of renewables together with P2X fuels offers energy security to many countries that are currently dependent on others. Finally, for the 13% of the word's population that still does not have access to reliable electricity, bringing cheap renewable energy to low income countries can accelerate energy inclusion and contribute to lifting around 1 billion people out of energy poverty.

The belief that the green energy transition will be expensive has been a major driver of the dismal response to climate change for the last forty years. However, this pessimism is at odds with past technological cost-improvement trends and also risks locking humanity into a much more expensive and dangerous energy future. Updating expectations to better align with historical evidence could fundamentally change the debate about climate policy and dramatically boost the likelihood of achieving a cheaper, greener and safer world.

References

Alberth, S. (2008), 'Forecasting technology costs via the experience curve—myth or magic?', *Technological Forecasting and Social Change* **75**(7), 952–983.

Anandarajah, G., McDowall, W. & Ekins, P. (2013), 'Decarbonising road transport with hydrogen and electricity: Long term global technology learning scenarios', *International Journal of Hydrogen Energy* **38**(8), 3419–3432.

Auffhammer, M. (2018), 'Quantifying economic damages from climate change', *Journal of Economic Perspectives* **32**(4), 33–52.

Brynolf, S., Taljegard, M., Grahn, M. & Hansson, J. (2018), 'Electrofuels for the transport sector: A review of production costs', *Renewable and Sustainable Energy Reviews* **81**, 1887 – 1905.
URL: *http://www.sciencedirect.com/science/article/pii/S1364032117309358*

Cai, Y., Judd, K. L., Lenton, T. M., Lontzek, T. S. & Narita, D. (2015), 'Environmental tipping points significantly affect the cost- benefit assessment of climate policies', *Proceedings of the National Academy of Sciences* **112**(15), 4606–4611.

Creutzig, F., Agoston, P., Goldschmidt, J. C., Luderer, G., Nemet, G. & Pietzcker, R. C. (2017), 'The underestimated potential of solar energy to mitigate climate change', *Nature Energy* **2**(9), 17140.
URL: *https://doi.org/10.1038/nenergy.2017.140*

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., Clack, C. T. M., Cohen, A., Doig, S., Edmonds, J., Fennell, P., Field, C. B., Hannegan, B., Hodge, B.-M., Hoffert, M. I., Ingersoll, E., Jaramillo, P., Lackner, K. S., Mach, K. J., Mastrandrea, M., Ogden, J., Peterson, P. F., Sanchez, D. L., Sperling, D., Stagner, J., Trancik, J. E., Yang, C.-J. & Caldeira, K. (2018), 'Net-zero emissions energy systems', *Science* **360**(6396).
URL: *https://science.sciencemag.org/content/360/6396/eaas9793*

De Cian, E., Buhl, J., Carrara, S., Bevione, M., Monetti, S. & Berg, H. (2016), 'Learning in integrated assessment models and initiative based learning-an interdisciplinary approach', *PATHWAYS Deliverable* **3**.

De Stercke, S. (2014), Dynamics of energy systems: A useful perspective, IIASA interim report, IIASA, Laxenburg, Austria.
URL: *http://pure.iiasa.ac.at/id/eprint/11254/*

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.

DeCarolis, J. F., Hunter, K. & Sreepathi, S. (2012), 'The case for repeatable analysis with energy economy optimization models', *Energy Economics* **34**(6), 1845–1853.

Eyre, N. (2019), 'Energy efficiency in the energy transition', *European Council for an Energy Efficient Economy Summer Study Proceedings* pp. 247–254.

Farmer, J. D., Hepburn, C., Mealy, P. & Teytelboym, A. (2015), 'A third wave in the economics of climate change', *Environmental and Resource Economics* **62**(2), 329–357.

Farmer, J. D. & Lafond, F. (2016), 'How predictable is technological progress?', *Research Policy* **45**(3), 647–665.

Grubler, A. (2010), 'The costs of the french nuclear scale-up: A case of negative learning by doing', *Energy Policy* **38**(9), 5174 – 5188. Special Section on Carbon Emissions and Carbon Management in Cities with Regular Papers.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421510003526*

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. & 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.
URL: *https://doi.org/10.1038/s41560-018-0172-6*

Heuberger, C. F., Rubin, E. S., Staffell, I., Shah, N. & Mac Dowell, N. (2017), 'Power capacity expansion planning considering endogenous technology cost learning', *Applied Energy* **204**, 831–845.

IEA (2019), The Future of Hydrogen, Technical report, International Energy Agency.

IEA & IRENA (2017), 'Perspectives for the energy transition: Investment needs for a low-carbon energy system'.

IRENA (2018), 'Global energy transformation: A roadmap to 2050'.

Koomey, J. & Hultman, N. E. (2007), 'A reactor-level analysis of busbar costs for us nuclear plants, 1970–2005', *Energy Policy* **35**(11), 5630 – 5642.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421507002558*

Koomey, J., Hultman, N. E. & Grubler, A. (2017), 'A reply to "historical construction costs of global nuclear power reactors"', *Energy Policy* **102**, 640 – 643.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421516301549*

Kramer, G. J. & Haigh, M. (2009), 'No quick switch to low-carbon energy', *Nature* **462**(7273), 568.

Krey, V., Guo, F., Kolp, P., Zhou, W., Schaeffer, R., Awasthy, A., Bertram, C., de Boer, H.-S., Fragkos, P., Fujimori, S., He, C., Iyer, G., Keramidas, K., Köberle, A. C., Oshiro, K., Reis, L. A., Shoai-Tehrani, B., Vishwanathan, S., Capros, P., Drouet, L., Edmonds, J. E., Garg, A., Gernaat, D. E., Jiang, K., Kannavou, M., Kitous, A., Kriegler, E., Luderer, G., Mathur, R., Muratori, M., Sano, F. & van Vuuren, D. P. (2019), 'Looking under the hood: A comparison of techno-economic assumptions across national and global integrated assessment models', *Energy* **172**, 1254–1267.
URL: *http://www.sciencedirect.com/science/article/pii/S0360544218325039*

Lafond, F., Bailey, A. G., Bakker, J. D., Rebois, D., Zadourian, R., McSharry, P. & 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.

Lafond, F., Greenwald, D. & Farmer, J. D. (2020), 'Can stimulating demand drive costs down? world war ii as a natural experiment', *INET Working Paper* .

Lovins, A. B., Ürge-Vorsatz, D., Mundaca, L., Kammen, D. M. & Glassman, J. W. (2019), 'Recalibrating climate prospects', *Environmental Research Letters* **14**(12), 120201.

McCollum, D. L., Zhou, W., Bertram, C., De Boer, H.-S., Bosetti, V., Busch, S., Després, J., Drouet, L., Emmerling, J., Fay, M. et al. (2018), 'Energy investment needs for fulfilling the paris agreement and achieving the sustainable development goals', *Nature Energy* **3**(7), 589.

McNerney, J., Farmer, J. D. & Trancik, J. E. (2011), 'Historical costs of coal-fired electricity and implications for the future', *Energy Policy* **39**(6), 3042–3054.

Mitchell, D., Heaviside, C., Vardoulakis, S., Huntingford, C., Masato, G., Guillod, B. P., Frumhoff, P., Bowery, A., Wallom, D. & Allen, M. (2016), 'Attributing human mortality during extreme heat waves to anthropogenic climate change', *Environmental Research Letters* **11**(7), 074006.

Mitsubishi Hitachi Power Systems (2019), 'Hydrogen power generation handbook'.

Nagy, B., Farmer, J. D., Bui, Q. M. & Trancik, J. E. (2013), 'Statistical basis for predicting technological progress', *PloS one* **8**(2), e52669.

Nordhaus, W. D. (2017), 'Revisiting the social cost of carbon', *Proceedings of the National Academy of Sciences* **114**(7), 1518–1523.

Pfeiffer, A., Hepburn, C., Vogt-Schilb, A. & 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), 054019.

Pindyck, R. S. (1999), 'The long-run evolution of energy prices', *The Energy Journal* **20**(2), 1–27.
URL: *http://www.jstor.org/stable/41322828*

Rahmstorf, S. & Coumou, D. (2011), 'Increase of extreme events in a warming world', *Proceedings of the National Academy of Sciences* **108**(44), 17905–17909.

Ram, M., Bogdanov, D., Aghahosseini, A., Oyewo, S., Gulagi, A., Child, M., Fell, H. & Breyer, C. (2019), 'Global energy system based on 100% renewable energy-power sector. study by lappeenranta university of technology and energy watch group, lappeenranta and berlin'.

Rentschler, J. E. (2013), 'Oil price volatility, economic growth and the hedging role of renewable energy', *World Bank Policy Research Working Paper 6603* .

Ricke, K., Drouet, L., Caldeira, K. & Tavoni, M. (2018), 'Country-level social cost of carbon', *Nature Climate Change* **8**(10), 895–900.

Staffell, I., Scamman, D., Abad, A. V., Balcombe, P., Dodds, P. E., Ekins, P., Shah, N. & Ward, K. R. (2019), 'The role of hydrogen and fuel cells in the global energy system', *Energy & Environmental Science* **12**(2), 463–491.

Stern, N. (2007), *The economics of climate change: the Stern review*, Cambridge University Press.

Trancik, J. E., Jean, J., Kavlak, G., Klemun, M. M., Edwards, M. R., McNerney, J., Miotti, M., Brown, P. R., Mueller, J. M. & Needell, Z. A. (2015), Technology improvement and emissions reductions as mutually reinforcing efforts: Observations from the global development of solar and wind energy, Technical report, MIT.

Trutnevyte, E. (2016), 'Does cost optimization approximate the real-world energy transition?', *Energy* **106**, 182–193.

Supplementary Information for Estimating the costs of energy transition scenarios using probabilistic forecasting methods

Rupert Way, Penny Mealy, J. Doyne Farmer

Correspondence to: rupert.way@smithschool.ox.ac.uk

Contents

List of symbols and abbreviations

Abbreviations

NPC	Net present cost
ICE	Internal combustion engine
EV	Electric vehicle
Li-ion	Lithium-ion
LCOE	Levelized cost of electricity
GHG	Greenhouse gas
P2X or PtX	Power-to-X
PtG	Power-to-gas
PtL	Power-to-liquids
CCS	Carbon capture and storage
EOR	Enhanced oil recovery

BAU	Business-as-usual
LHV	Lower heating value
HHV	Higher heating value
CCGT	Combined-cycle gas turbine
PV	Photovoltaic
VRF	Vanadium redox flow
COP	Coefficient of performance
PPA	Power purchase agreement
bbl	Barrel of crude oil
MMBtu	Million British thermal units
DSM	Demand-side management
SMR	Steam methane reforming
CRF	Capital recovery factor
AEL	Alkaline electrolysis
PEM	Polymer-electrolyte membrane
SOEC	Solid oxide electrolyzer cell
SOFC	Solid oxide fuel cell
VRE	Variable renewable energy

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^k_{slack}	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
$EF_t^{\mathrm{j},\ total}$	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*) forecast net energy system 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 historical data to generate empirically grounded probability distributions of energy system costs, and calculate the expected net present costs of different scenarios. This reveals the preferred scenario in terms of expected costs. Our perspective is that the future is intrinsically uncertain and decision making is always a gamble. We thus reformulate the energy transition question in terms of probabilities, essentially viewing it as making a bet on future technological progress, conditional upon observed past progress. We answer the question: given the consistent cost reductions in several clean energy technologies observed over the last three decades, what is the rough shape (in terms of time-scales and energy mix) of a transition to a future energy system that is likely to be cheapest? Note that our model only considers costs directly associated with supplying energy services to a growing global economy; we do 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 calculate the expected net present cost (NPC) of three transition scenarios to a clean energy system. The model only considers the very largest energy flows and most important technology substitutions, and ignores everything that, from our current perspective, appears negligible. We model the core components of the global energy system each year from 2019 to 2070 using 13 energy supply technologies: 3 fossil fuels, 7 electricity generators and 3 energy storage technologies. The model is global and does not differentiate among different geographic regions. The benchmark scenario is called *No Transition*, and the three others are called *Fast Transition*, *Slow Transition* and *Slow Nuclear Transition*. These are shown in detail in Section 4, but we also refer to them repeatedly before that in order to explain the scenario construction process. 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 total relative costs.

There are two steps to the overall modeling procedure. First, the four energy scenarios are constructed, subject to the constraint that they must all provide identical levels of total energy service supply. There are no market mechanisms involved in the construction process: scenarios are created by specifying exogenously varying technology growth rates. Second, technology cost forecasting methods are used to forecast individual technology costs in each scenario. These are discounted and summed to give probability distributions of the NPC of each scenario, and to compute the expected NPC of transition for each scenario, relative to the No Transition scenario.

As we discuss further in section 1.6, this simple two-step approach contrasts with most other global energy system models, which usually involve modeling 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 closely connected to empirical data. The simplicity and transparency of the model and its results make it easy to be understood, replicated and modified with 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 qualitatively influence the results (see Table S1 for a breakdown). This allows us to focus on a few critical technology substitutions that can have a major impact on the energy system. We also aim to only include key features of the energy transition that differ dramatically between scenarios. For example, when modeling the transition from internal combustion engine (ICE) vehicles to electric vehicles (EVs), we ignore the costs of capital stock replacement because these costs are expected to be roughly equivalent[1]. We do, however, model the cost of providing the energy required to power the fleet and provide the basic mobility service in each scenario (i.e. the cost of gasoline versus the cost of electricity).

Similarly, regarding the costs of infrastructure conversion and expansion, we consider the requirements in different scenarios, and ensure that key technologies are provided in sufficient quantities so that any residual conversion and expansion costs are likely similar across all scenarios, and can therefore be neglected. For example, in a very high renewables scenario, *either* significant power network expansion *or* significant energy storage is required (or some combination). Our approach is to ensure that large renewables deployments are always accompanied by large quantities of energy storage technologies, so then we can assume that any resid-

[1] For example, BNEF's EV Outlook 2019 (BNEF 2019) says "we expect price parity between EVs and internal combustion vehicles (ICE) by the mid-2020s in most segments, though there is wide variation between geographies and vehicle segments"

ual network expansion requirements will be similar across scenarios. This is reasonable when one considers that even in a very low renewables scenario significant network upgrades and expansion would be required, as outdated infrastructure is retired, and new systems are installed to meet the increasing global electricity demand.

Regarding fuels, we argue that transportation, distribution and storage costs are approximately equivalent across all scenarios, so can also be neglected. Undoubtedly there will be many specific technologies where conversion costs are a concern and the transition will not be smooth (residential heating in old properties for example), but we argue that these contribute only a small fraction towards the total system costs, so should be omitted in order to better capture the most important transition dynamics.

1.2 Energy system description

We describe the global energy system in standard terms (see GEA (2012) 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 any biomass-derived fuels. It is not very clear when used to describe 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.

- *End-use applications* are the last step in the energy delivery chain, where energy gets used by a customer or end-user to perform some useful energy service.

- *Energy carriers* are substances or phenomena used in end-use applications, e.g. coal, electricity, power-to-X (P2X) fuels. Primary energy resources are often converted to different energy carriers to facilitate transport, storage or market transactions of energy.

- *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 well defined though, and it is applied flexibly in different contexts. In this model we use a very specific definition, described below.

- *Final energy* is the energy delivered to energy users, either in end-use sectors of the economy, or in the energy sector itself. It is the energy embodied in energy carriers as they arrive at energy users for consumption via various energy technologies.

- *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 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. 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 EV or an ICE vehicle. 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 ICE vehicle 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* **in this model** (also referred to as simply *technologies*) are defined to be precisely those components of the model whose costs are specified and included in the total system cost calculation. This choice is motivated by the fact that most model components' costs are obtained via a *technology cost forecasting model*.[2]

[2]This definition coincides with the familiar understanding of energy technologies in most cases, though there is one exception. In end-use applications (e.g. oil in transport, coal in industry) the commodities oil, coal and gas are all considered to be "energy

While we note that the concept of primary energy is frequently used by energy agencies and many energy companies, 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

Since the goal of the work here is to calculate the costs associated with energy transition scenarios, a critical bottleneck is the availability of cost or price data for various energy technologies. The model is constructed so as to draw on all available historical data, and hence capture long-term trends in technological progress as accurately as possible. At the same time, however, the model must also be sufficiently technically realistic.

For crude oil, coal, natural gas, and coal electricity, consistent annual price data exists for around a century. For nuclear, solar and wind electricity, both energy costs and capacity costs are available for a few decades starting between the 1970s and 90s. Cost data for other electricity generation technologies are only available since the 1990s. Battery and electrolyzer costs are only available for a few decades. For many other energy technologies only very limited or irregular data exists, making it hard to construct satisfactory cost forecasts.

The model components and overall strategy are designed to make the best use of available data, while also representing the global energy system with as much accuracy as possible. The key datasets 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 PV (medium or short time series).
- Capital cost of: Li-ion batteries ($/kWh), Vanadium redox flow batteries ($/kWh) and electrolyzers ($/kW) (mostly short, irregular time series).

Data and variable choices are discussed further in Section 5. Based on these considerations, 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.
- 3 direct-use primary energy resources: 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. (Note that daily batteries are used in the intermediate and transport sectors separately.)

The model system is shown in Figure S1. End-use sectors supply useful energy to provide energy services to consumers; energy carriers provide final energy to end-use sectors and the energy sector for conversion to useful energy; and energy technologies provide the energy, storage and conversion capabilities required to produce and manage the energy carriers. 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.

Current costs of all these energy technologies are known, and their future costs may be forecast based on relatively good historical data. The energy system components included in the model currently account for approximately 97% of electricity generation and 95% of CO_2 emissions from the power sector; and 83% of all final energy from end-use sectors and 82% of final energy CO_2 emissions. Due to our assumptions about traditional biomass though (described below), the model implicitly accounts for a further 7% of current final energy and 11% of final energy CO_2 emissions, taking the totals to 90% and 93% respectively.

Note that we focus entirely on energy supply technologies in this paper. There are many possibilities for innovation, efficiency increases and improvements in demand-side technologies that could further lower

technologies" in their own right, but P2X fuels are not. This is because the costs of end-use oil, coal and gas are accounted for and included in the total system cost calculation, but P2X fuel production costs are only accounted for *indirectly*, as the sum of its constituents' costs: electricity generation and installed electrolyzer capacity.

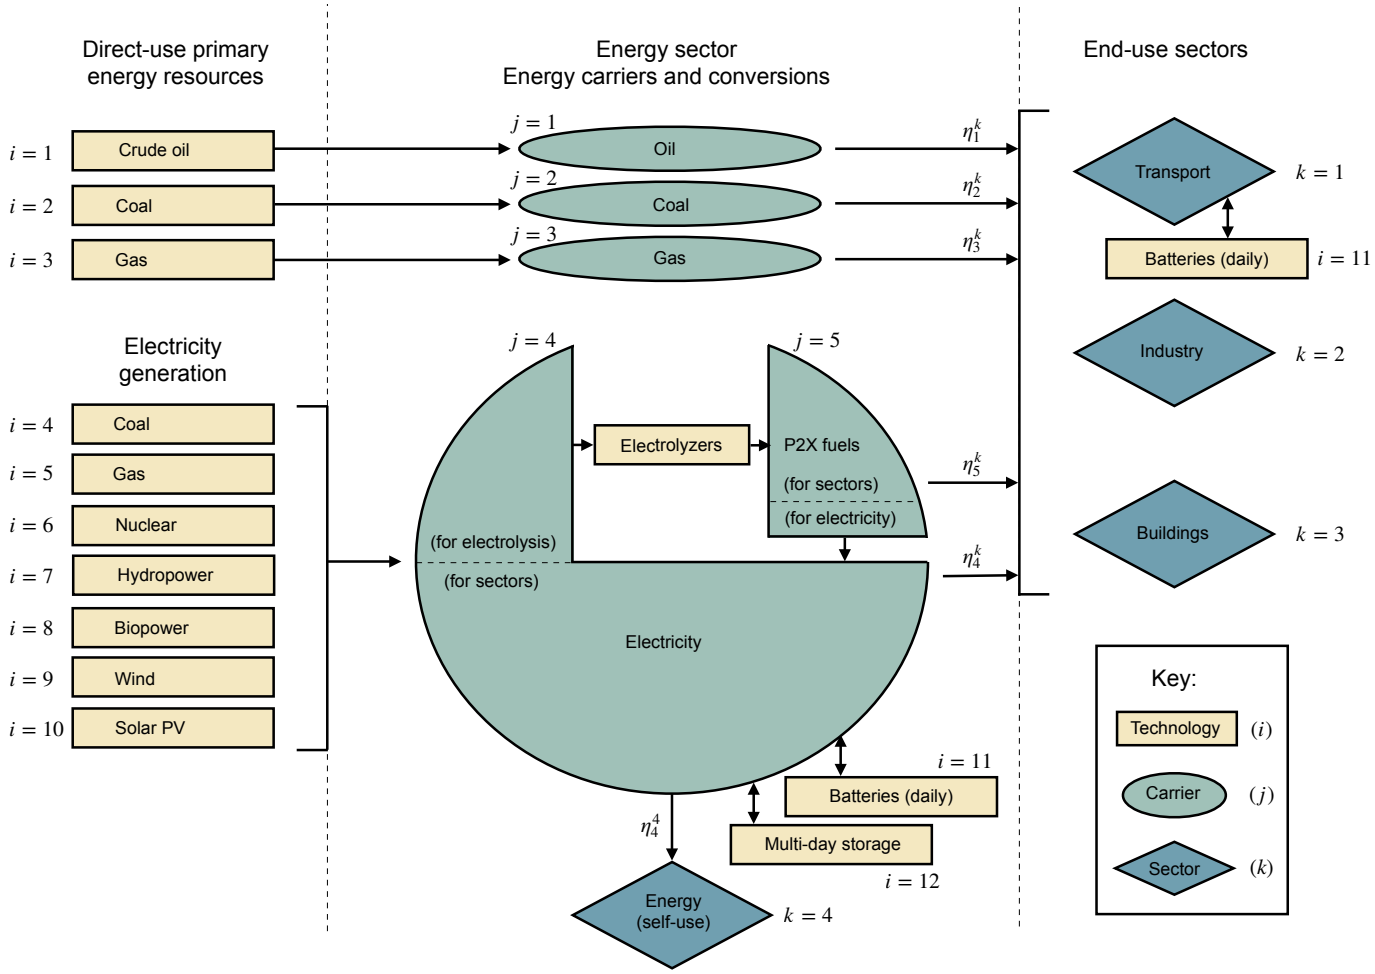

Figure S1: Energy system model components. Historical costs are known for all technologies, and their future costs are modeled. The system cost is then approximated by summing up forecast costs of all technologies, in suitable quantities to meet end-use sector requirements. Indices and conversion efficiencies are described in Sections 2.4 and 1.5.

the costs of transition scenarios considered here (Wilson et al. 2012, Creutzig et al. 2018, Lovins 2018). We exclude such considerations on the grounds of simplicity and communicability, since they are not required to demonstrate out main result.

Table S1 shows which electricity generation technologies are included in the model and which are omitted. Table S2 gives a detailed breakdown of final energy in end-use sectors and other major uses, and shows which energy flows are included in the model.

	Electricity generation, EJ	
Technology	Included	Excluded
Oil		2.9
Coal	36.4	
Gas	22.0	
Nuclear	9.8	
Hydro	15.1	
Biopower	2.3	
Wind	4.6	
Geothermal		0.32
Solar PV	2.1	
CSP		0.043
Marine		0.004
Totals	92.4	3.3
	96.6%	3.4%

Table S1: Electricity generation in 2018 by technology, as listed in the IEA WEO 2019, showing whether the technology is included or excluded from our model. (Column sums may differ from totals due to rounding.)

Energy carrier	End-use sectors: Transport, EJ Incl.	Transport, EJ Excl.	Industry, EJ Incl.	Industry, EJ Excl.	Buildings, EJ Incl.	Buildings, EJ Excl.	Intermediate/other uses: Energy, EJ Incl.	Energy, EJ Excl.	Other, EJ Incl.	Other, EJ Excl.	Totals by carrier, EJ Incl.	Totals by carrier, EJ Excl.
Oil	110.1		12.5		13.8						136.4	
Coal			33.3		5.23						38.5	
Gas			27.0		29.3						56.3	
Electricity	1.3		33.6		42.33		15.5				92.8	
Heat				5.99		6.28						12.3
Bioenergy (excl. Trad. Biom.)				8.92		4.9						13.8
Traditional biomass						26.0						26.0
Biofuels		3.7										3.7
Other fuels		4.7										4.7
Other renewables				0.04		2.0						2.0
Petrochem. feedstock										23.0		23.0
Unspecified								47.5		22.7		70.3
A. *Entire system*:												
Column totals, EJ	111.4	8.4	106.4	15.0	90.7	39.2	15.5	47.5	0.0	45.7	324.0	155.7
Percentage of sector	93%	7%	88%	12%	70%	30%	25%	75%	0%	100%	68%	32%
Sector totals, EJ	120		121		130		63		46		480	
B. *Reduced system*: (excluding Heat, Trad. biomass, Petrochem. feedstock, & Unspecified)												
Column totals, EJ	111.4	8.5	106.4	8.9	90.7	6.9	15.5	0.0	0.0	0.0	324.0	24.3
Percentage of sector	93%	7%	92%	8%	93%	7%	100%	0%	-	-	93%	7%
Sector totals, EJ	120		115		98		16		0		348	
Summary coverage of reduced system	93%		92%		93%		100%		100%		93%	

Table S2: Final energy and other major uses in 2018 by sector and energy carrier, as listed in the IEA WEO 2019, showing whether each component is included or excluded from our model. Each column for a given sector is split into two columns, one showing what is included and one showing what is excluded. (Row and column sums may differ from totals due to rounding.) Row block A shows the model coverage relative to the entire energy system, including out-of-scope components such as heat, traditional biomass and petrochemical feedstock. Row block B shows the model coverage relative to the reduced system that results from excluding these out-of-scope variables.

1.3.1 Oil, coal, gas and electricity

In line with our approach of only considering the largest energy flows, the only components of the current energy system that we model are oil, coal, gas and electricity. There are many other energy carriers, but we exclude these from the model because of their relatively small impact on the system.

Note that we treat coal and gas used directly in end-use sectors ($i = 2, 3$) as different technologies from coal and gas electricity generation ($i = 4, 5$). This is because all electricity generation technologies are considered to be substitutes (subject to sufficient energy storage provisoin), but substituting direct-use fuels for other energy carriers requires accounting for their different conversion efficiencies.

1.3.2 Heat

Heat is omitted from the model because it makes up only a small fraction of the system and is qualitatively different than the other energy carriers and technologies we use. 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 in the same modeling framework as the other system components.

1.3.3 Traditional Biomass

The one substantial end-use energy flow we do not model is traditional biomass used in buildings, which currently makes up 20% of final energy used in the building sector (26 EJ of 130 EJ). This refers to fuelwood, charcoal, animal dung and agricultural residues used for cooking and heating. We do not model this explicitly because it is assumed that the annual 2% expansion in global useful energy supply will gradually subsume this category of energy use, as modern, efficient technologies become more widespread. Furthermore, it is a carbon neutral energy flow, so does not need to be accounted for in the way that other energy sources do (at least from an emissions perspective).

1.3.4 Petrochemical feedstock

Petrochemical feedstock is omitted from the model. This 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 category from our model on the basis that these fossil fuels are not used as energy carriers in the usual 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 model (in the industry sector), it is just the feedstock itself that is not. Thus our model covers the emissions resulting from the feedstock processing, so that whether or not the feedstock itself is composed of fossil fuels is irrelevant to our analysis (providing any substitute feedstock requires similar levels of energy for processing, of course).

1.3.5 Energy sector

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", see the Definitions section in Annex C of IEA (2019*c*). We include the electricity consumption of the energy sector (15.5 EJ/yr) in the model but exclude its fossil fuel consumption (47.5 EJ/yr) (labelled "Unspecified" in the table). This is because the future development of the energy sector will depend strongly on which scenario is pursued (it is the only such sector). Its current consumption of nearly 50 EJ of fossil fuel energy per year would decrease towards zero in a Fast Transition scenario and, crucially, would not need to be replaced, it would simply disappear. Our model is not designed for such scenario-dependent sector growth (or decline), and cannot deal with this behaviour without the introduction of further model complexity. Instead we keep things simple and just exclude the fossil fuel component completely, while assuming that the electricity component grows as usual.[3] Despite omitting these fossil fuels from the energy balance in the model we do in fact include a portion of the associated costs, by adding a constant oil refining cost to each EJ of oil used in end-use sectors (see Section 6.1). There are many other transformation processes besides oil refining though, whose costs this very simple method omits, so it likely underestimates the total fossil fuel costs significantly. Thus we are being highly conservative in our fossil fuel cost assumptions, for the sake of model simplicity.

1.3.6 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 would have very high environmental costs. 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 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.

[3] It is essential to include the electricity component in the model because, first, electricity technology experience curves rely on accurate generation quantities, and second, we must ensure that enough storage capacity exists to cope with all nondispatchable electricity used, including that used in the intermediate energy sector.

1.3.7 CCS

Carbon capture and storage (CCS) technologies are omitted from our model because their production scale so far is very small and their commercial viability is unclear. CCS costs are generally not well known, as this is sensitive commercial information for fossil fuel companies (who are the only organizations active in the space). Carbon capture for enhanced oil recovery (EOR) is a well established technology that has been operating commercially in many locations around the world since 1972. Despite decades of investment in R&D programs in both the public and private sectors, there is no historical evidence that costs have fallen, and no evidence that they will fall in future to levels allowing CCS to contribute seriously to emissions mitigation without a high carbon price. In fact costs have increased since 2005 (Rubin, Davison & Herzog 2015). Historical analysis suggests that bringing a technology from the research phase to production typically requires decades (Grubler et al. 1999), so it seems unlikely to play a decisive role on the timescale required.

We distinguish between CCS whose purpose is to burn fossil fuels without net emissions, and CCS whose purpose is to remove existing carbon from the atmosphere. It may be necessary in the future to do the latter in order to offset the emission of greenhouse gases (GHG) in agriculture, or to lower GHG levels below the levels that we will arrive at once the green energy transition is complete. We do not model this.

1.3.8 P2X fuels

P2X fuels here refer to several fuels, including hydrogen, ammonia and methane. 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 gasification (IEA 2019*b*). Hydrogen produced from fossil fuels with unconstrained GHG emissions is known as "grey hydrogen" and accounts for almost all hydrogen ever produced (Schoots et al. 2008). Fossil fuel hydrogen with CCS is known as "blue hydrogen", though this is rarely produced due to the extra costs incurred by CCS. (Despite this, SMR with CCS is currently believed by some to be a viable option in future.)

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 feasible in the long term. Second, blue hydrogen is dependent on CCS technologies, which 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 Section 6), so there is little reason to expect any significant reduction in this component of the resulting hydrogen cost. Finally, both SMR and gasification are mature technologies, so there is little reason to expect any significant reduction in this cost component either. Unlike green hydrogen, neither grey nor blue hydrogen are expected to benefit from any future cost reductions in relatively immature, fast progressing technologies such as solar and wind electricity, batteries, or electrolyzers. We therefore exclude them from the model.

We do not attempt to model the final mix of P2X fuels as this would be too complicated and does not qualitatively affect the results. Power-to-hydrogen represents a majority of the cost for all P2X fuels, and we only model this step. For example, producing ammonia via hydrogen currently costs about 10-25% more than the cost of producing hydrogen alone (IEA 2019*b*), while the premium for synthetic methane is around 30% (Brynolf et al. 2018).

1.3.9 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.). However, just as we do not model the exact P2X fuel mix, we also do not model the exact infrastructure requirements.

Hydrogen infrastructure currently costs around 2-3 times as much as oil and gas infrastructure per unit of energy (see Section 3). This is partly because hydrogen gas has very low energy density (3-4 times lower than natural gas), so more compression is required during transmission and storage, but also partly due to the lack of scale of the network. There are currently about 5000 km of hydrogen pipelines around the world, and 3 million km of natural gas pipelines. However, many 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. Infrastructure costs are generally not believed to be a significant impediment to deployment (it is the cost of the hydrogen itself that currently prevents more widespread use, at 3-4 times the cost of natural gas). Ammonia is cheaper and much more convenient to store and distribute than hydrogen, due to its higher density, and over 7000 km of ammonia pipelines currently exist. Indeed, 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 (IEA 2019*b*)).

Due to lack of data on infrastructure costs for pipelines and storage through time, 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. As it turns out, in the Fast Transition scenario the size of the required P2X industry is 2-3 times smaller than the size of the combined oil and gas industries in alternative scenarios. This is because most residential uses are electrified, and the fuel-based energy distribution and storage network consolidates around industrial uses, shipping and aviation. This results in a 2-3 downscaling of the network. 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. In fact this is a conservative assumption because P2X infrastructure costs are likely to decrease with scale. Hence total P2X infrastructure costs in the Fast Transition scenario will 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 comes to scale.

As noted earlier, we also assume electricity infrastructure costs are approximately equivalent across scenarios, because high levels of renewable electricity generation are always matched by high levels of electricity storage, whose costs are included in the system cost calculation. Thus we consider all infrastructure requirements to be part of the routine capital stock replacement and improvement process that takes place over several decades in all scenarios. We discuss these issues in more detail in Section 3.

1.4 Variables and units

In order to keep our model as simple as possible, we deal almost entirely in terms of energy, as opposed to power. While we note that energy technologies have multiple characteristics that interact in complex ways, this approach allows us to calibrate the model and produce cost forecasts in a simple, transparent way given the available data.

Many of the data sources used in this model are described in terms of different units. To make all quantities and costs comparable, we have converted all data to exojoules (EJ) and dollars per gigajoule ($/GJ; 1 $/GJ = 1 billion$/EJ). The only exception to this energy-only perspective is electrolyzers, which are described in terms of installed capacity; data sources use kW and $/kW but we use EJ hr^{-1} and $/GJ hr^{-1} (given an annual capacity factor assumption).

1.5 End-use conversion efficiencies

End-use 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 $\eta^{sector}_{carrier}$. 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 $\eta^{industry}_{gas}$). Evidence for the validity of this assumption, and the specific values we use, are given in De Stercke (2014). It is beyond the scope here to predict how these efficiencies may change in the future, so we assume 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 S3. Also shown are current final energy values and the useful energy values implied by the conversion efficiencies.[4]

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 ICE vehicles) 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

[4]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 84 MW of hydrogen electrolyzers (of the appropriate type) currently exist. Assuming a capacity factor of 50% yields a current annual quantity of P2X fuels of 13.25 $\times 10^{-4}$ EJ/yr. For simplicity we suppose that this is split equally among the three end-use sectors, giving values of 4.4 $\times 10^{-4}$ 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 Section 6.12 for further details.

Sector	Carrier	Final → Useful efficiency, $\eta^{sector}_{carrier}$	Current final energy, Mtoe	Current final energy, EJ	Current useful energy, EJ
Transport	Oil	0.25	2629	110.1	27.5
	Electricity	0.8	32	1.34	1.07
	P2X fuels	0.5	-	4.4×10^{-4}	2.2×10^{-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	-	4.4×10^{-4}	2.6×10^{-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	-	4.4×10^{-4}	2.6×10^{-4}
	All		2166	90.6	72.7
Energy	Electricity	1.0	371	15.5	15.5
	All		371	15.5	15.5

Table S3: Carrier and sector specific conversion efficiency assumptions. $\eta^{sector}_{carrier}$ is the final to useful energy conversion efficiency of the given carrier in the given sector. These values are chosen based on the data and analysis presented in De Stercke (2014). Final energy data for 2018 is shown (from the World Energy Outlook (IEA 2019*c*)), along with the useful energy values implied by the assumed conversion efficiencies.

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 chosen for the use of electricity in buildings. This was based on historical data, but may vary considerably in a future where heating is mostly provided by heat pumps. These devices use ambient energy drawn from the environment (unaccounted for in all energy models, including this one), and have a coefficient of performance (COP) (i.e. conversion efficiency) in the range 2-5 (Staffell et al. 2012). If heat pumps become widespread, the average conversion efficiency in the sector could conceivably rise to $\eta^{buildings}_{electricity} \approx 2$ or higher, which would significantly reduce the quantity of final energy required by the sector. Thus our value of 1 here is conservative. No other carrier-sector pairs suffer from this large uncertainty.

1.6 Comparison to existing energy system modeling approaches

There are a plethora of energy system models that have been designed to better manage 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), modeling techniques, sectoral and geographical coverage and technological detail (see Bhattacharyya & Timilsina (2010) for a review). Many models (e.g. MESSAGE (Messner & Strubegger 1995), MARKAL/TIMES (Loulou et al. 2004, 2005), OseMOSYS (Howells et al. 2011)) 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) (DeCarolis et al. 2017). However, there are numerous other models employing different modeling approaches, such as POLES (Keramidas et al. 2017), which is based on a partial-equilibrium, dynamic-recursive simulation and LEAP (Heaps 2008), which employs an accounting-based framework.

The model developed in this paper departs from traditional conventions in energy system modeling 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" (DeCarolis et al. 2017, pp. 185), we have deliberately aimed to develop a novel energy system modeling approach that is as parsimonious, transparent and empirically grounded as possible. We consequently only represent energy system components that are necessary to model 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 BAU*: most traditional energy system models construct a BAU 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). However, merely constructing a BAU scenario in this way requires many strong assumptions about population growth, economic growth, technology costs, technological progress and innovation, and the role of fossil fuels in society. 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 trajectory (Ellenbeck & Lilliestam 2019). 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 scenarios (Trutnevyte 2016). As our model is designed to explore objectively defensible future energy system transition scenarios that are empirically grounded in historical technological trends, we instead construct a BAU on the basis of a trajectory in which almost no technological change occurs within the energy system, such that the current technological structure continues into the future. Our BAU scenario is just a continuation of the constant-growth, near-static-structure energy system that has existed for at least the last fifty years (as a result of innumerable policy decisions). We then compare this constant growth "No Transition" BAU to alternative scenarios that *do* involve changes in the underlying technological structure of the global energy system.

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 (Berglund & Söderholm 2006, Anandarajah & McDowall 2015). 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 costly (DeCarolis et al. 2017). While more recent efforts have reduced computing time by employing data clustering approaches (Heuberger et al. 2017), 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 cheaper (Ellenbeck & Lilliestam 2019). 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 limits (Loulou & Labriet 2008).

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 transitions (Trutnevyte 2016)), we exogenously construct scenarios using a few simple rules constraining the supply, growth and substitutability of technologies, then use transparent, empirically 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 (Way et al. 2019). We instead explore a few representative areas of the space of scenarios, and demonstrate—without the imposition of ad-hoc assumptions or unnecessary model complexity—that rapidly transitioning to a global energy system based on renewables will likely generate substantial economic benefits – even without accounting for the avoided costs of climate damages.

2 Scenario construction

We follow Grubler et al. (2018) 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 will approximately 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 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 we:

1. specify useful energy provided by each sector (exogenous assumption, identical for all scenarios),
2. specify energy carrier quantities used in each sector (scenario-specific),
3. specify proportion of nondispatchable generation in electricity mix in 2100 (scenario-specific),
4. increase P2X fuels production to account for intermittency if required (dependent on 3),
5. increase total electricity generation to account for electrolytic production of P2X fuels,
6. specify electricity generation technology quantities (scenario-specific),
7. calculate quantity of daily-cycling batteries (a function of nondispatchable generation and electrification of transport sector),
8. calculate quantity of multi-day storage (a function of nondispatchable generation quantity),
9. calculate electrolyzer capacity required to produce total quantity of P2X fuels.

We use a time horizon of 2070 for our main results, but also present robustness checks on different time horizons up to 2090 in Sections 7.2 and 7.3. 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

In general we construct scenarios by growing or contracting different model components at different rates over time. To implement these growth processes a common framework is used: growth rates follow logistic growth curves (i.e. S-curves). Each model component begins in an initial state corresponding to its current production level and roughly matching its historical growth rate.[5] Then over time each component gradually transitions to a new configuration, ending up in a new steady state defined by its long-term growth rate.

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

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

Each model 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 $t_0 + t_1$. From $t_0 + t_1$ to $t_0 + t_2$ 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. For simplicity, a single steepness parameter coefficient equal to 50 is used for all model components.

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

$$q_{t+1} = (1 + g_t)q_t, \qquad t = 1, 2, \ldots, T, \;\; g_t > -1. \tag{2}$$

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.

[5] This is the general methodology but there are a few exceptions. For example, in order to implement the No Transition scenario, the growth rate of solar PV must fall sharply very quickly; to produce this effect we just use a very low initial low growth rate.

(*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

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

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 $g_T = 0.02$.

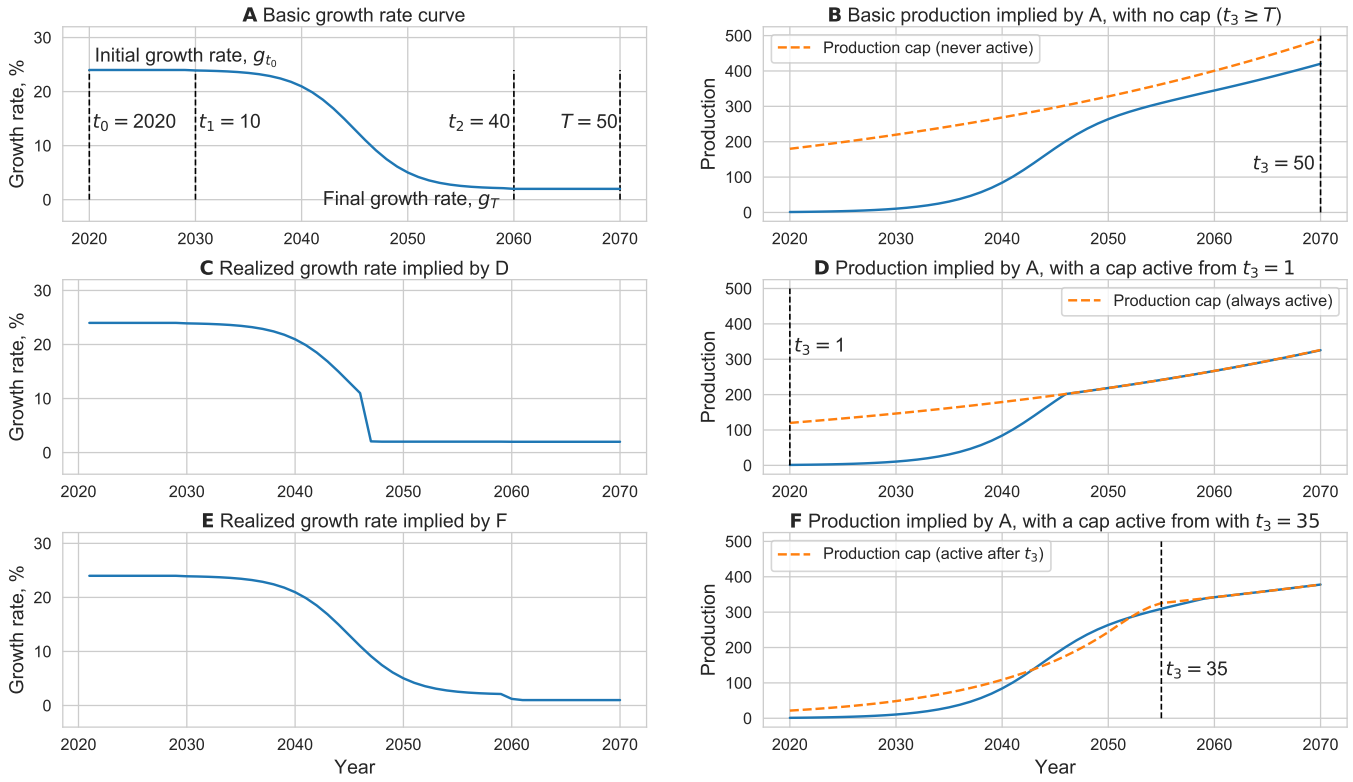

Figure S2: Growth processes used in scenario construction. Panel *A* shows a growth rate transitioning from 24% to 2% per year following a logistic growth curve. The growth rate g_t is defined by Eq. 1. This corresponds to a huge expansion in production over 20 years, followed by a gradual slowdown to a new long-term steady state, as shown in Panel *B* (case (*i*), 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 realized 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 realized growth rate of production in this case.

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 average (IEA 2019*a*). 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) (Ritchie & Roser 2018). 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.[6]

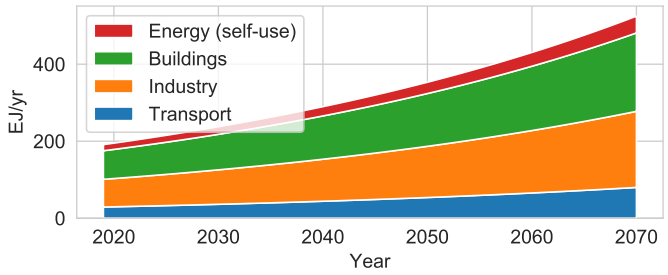

Figure S3: Exogenous assumption of useful energy demand growth.

We require that the quantity of useful energy provided by each sector is exactly the same in all scenarios, so that identical levels of energy service are supplied. In contrast to some other energy models, in particular partial equilibrium models, in which energy service consumption may vary depending on energy prices, this allows for simple, transparent and meaningful (apples-to-apples) comparison of the costs associated with different scenarios.

2.3 Indices, slack variables and meeting the useful energy constraint

With the total useful energy supplied to each sector completely determined in this way, we first grow and contract the different energy carriers to meet the totals in each year, then grow and contract the direct-use primary energy 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 grow at specified rates, then the various energy carriers are grown or contracted to meet the end-use demands, and in turn various technologies are grown and contracted to meet the carrier demands.

In each case, in order to meet a specific constraint, all input components are grown or contracted using the method in Section 2.1 except for one: this "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. They are used to supply any residual demand left unfulfilled by the sum of all other components (for more detail see Eq. 7 below).

Any mix of primary energy and electricity generation resources that meets the carrier constraints in each year (and hence also the useful energy constraints) is a valid scenario. The costs of such valid scenarios may be compared fairly with one another, as they are guaranteed to 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 $N \times T$. Here $N = 13$ and $T = 52$, so each portfolio has dimension 676. This space is far too large to explore thoroughly, 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 $t_0 = 2018$ 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 primary energy or electricity generation technologies used as input to carrier j. Values are shown in Table S4.

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 S3). Let $EU_t^{j,k}$ be the associated useful energy at time t. As described in Section 1.5, let the conversion efficiency of final to useful energy of carrier j when used in sector k be η_j^k. Then, by definition,

$$EU_t^{j,k} = \eta_j^k \, EF_t^{j,k}. \tag{4}$$

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

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	
j	1	Carrier	Oil	$\mathcal{I}_1 = \{1\}$
	2	Carrier	Coal	$\mathcal{I}_2 = \{2\}$
	3	Carrier	Gas	$\mathcal{I}_3 = \{3\}$
	4	Carrier	Electricity	$\mathcal{I}_4 = \{4, 5, 6, 7, 8, 9, 10\}$
	5	Carrier	P2X fuels	
k	1	Sector	Transport	$\mathcal{J}_1 = \{1, 4, 5\}$
	2	Sector	Industry	$\mathcal{J}_2 = \{1, 2, 3, 4, 5\}$
	3	Sector	Buildings	$\mathcal{J}_3 = \{1, 2, 3, 4, 5\}$
	4	Sector	Energy (self-use)	$\mathcal{J}_4 = \{4\}$

Table S4: 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 ($\mathcal{I}_k$). See Figure S1 also.

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:

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

The RHS values are known at time t_0, as shown in Table S3, and $EU_t^{total,k}$ is calculated for all times using the method in Section 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, $j \in \mathcal{J}_k \backslash \left\{j_{slack}^k\right\}$, 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.[7] 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:

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

and the final energy is calculated with Eq. 4.

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 zero carbon energy sources. Electrifying each sector as much as possible is believed to be necessary for a clean energy transition, 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.

[7] 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.

Regarding transport, Mai et al. (2018) 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 would 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 would still be rising, and that light-duty cars and trucks make up of the order of 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, though this is more speculative). This just leaves some heavy-duty trucking, plus long-distance air and sea transport. We therefore set $\xi^{4,1} = 0.8$, so that a maximum of 80% of useful energy in transport is delivered by electricity. This allows for the full electrification of light duty vehicles, but 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. (2018) 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. That said, most industrial processes are currently not designed to use electricity, and electrified substitutes are often not yet available, so barriers to adoption are currently high. 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". We therefore use a much more 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. (2018) 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): $\xi^{4,1} = 0.8$, $\xi^{4,2} = 0.75$, $\xi^{4,3} = 0.9$.[8] 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 (conditional upon any given energy carrier mix) will be lower.

Although it is certainly possible to represent all the varied circumstances of energy consumption in end-use sectors with much greater accuracy 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 modeler 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 and the sum of the useful energy provided by non-slack carriers:

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

We keep the total amount of energy provided by slack carrier variables small so that the precise assumptions about the slack variable do not affect the conclusions of the model. The quantity of final energy delivered by each carrier to sector k is then given by Eq. 4. By applying this procedure for each sector we construct the annual quantities, in EJ, of final and useful energy of all energy carriers required to meet the useful energy constraints.

[8] We also set $\xi^{4,4} = 1$ since this model only considers the energy supplied to the energy sector in the form of electricity, as previously noted — this "sector" is fully electrified already and remains so.

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 used in year t be $Elec_t$. This is the sum of electricity use in sectors and in P2X fuels production. Transmission and distribution losses are small (around 5% in OECD countries (Kassakian et al. 2011)) and falling, due to ever more advanced equipment; we ignore them in this analysis. The total electricity generated each year is the sum of generation from the individual technologies:

$$Elec_t = \sum_{i \in \mathcal{I}_4} q_t^i \tag{8}$$

The initial values at time t_0 are all known, as shown in Table S5.

Source	Electricity generation in 2018	
	TWh	EJ
Coal	10101	36.4
Gas	6183	22.3
Nuclear	2701	9.7
Hydroelectric	4193	15.1
Biopower	623	2.6
Wind	1270	4.6
Solar PV	585	2.1

Table S5: Initial electricity generation data, from BP 2019.

Next pick a slack electricity generation variable, denote its index by $i_{slack} \in \mathcal{I}_4$. For each non-slack technology, $i \in \mathcal{I}_4 \backslash \{i_{slack}\}$, 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. For example this occurs with nuclear power in the Slow Transition scenario, see Table S8, Figure S9. Nuclear begins with around 10% of the electricity mix and has a long-term upper limit of 12% of the generation mix. Initially it grows fast so that it breaks the 12% mix limit, but as other technologies also grow rapidly, nuclear's share eventually decreases, so that after t_3 it can continue growing slowly towards its final 12% share of the mix.

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

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

and the slack technology is subsequently given by

$$q_{t+1}^{i_{slack}} = Elec_{t+1} - \sum_{i \in \mathcal{I}_4 \backslash \{i_{slack}\}} q_{t+1}^i. \tag{10}$$

2.6 Direct-use fuels

The quantities of direct-use oil (i=1, j=1), coal (i=2, j=2) and gas (i=3, j=3) required in a scenario (i.e. their primary energy quantities) are simply their associated final energy quantities in the energy carrier calculations above:

$$q_t^1 = EF_t^{1,total} \tag{11}$$

$$q_t^2 = EF_t^{2,total} \tag{12}$$

$$q_t^3 = EF_t^{3,total} \tag{13}$$

We have now specified the energy quantity requirements for all sectors, energy carriers, and energy supply technologies (i = 1-10). The model components that remain to be quantified are the energy storage and conversion technologies ($i = 11, 12, 13$), which we address next.

3 How we solve the intermittency problem

Because solar power and wind power are intrinsically variable, as they become the dominant sources of power we must be able to store energy to deal with the intermittency problem. The best solution to this problem depends on the timescale. We model three different technological solutions on short, intermediate and long timescales. Short timescales are defined as being less than a day; intermediate timescales as one to three days, and long timescales as more than three days.

The key reference for the modeling decisions we make regarding storage and flexibility is the paper by Blanco & Faaij (2018). 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.

There are at least six solutions to cope with fluctuations in the supply and demand of electricity, including:

1. *Demand side management*, which means scheduling demand for electricity so that it occurs when there is available supply.

2. *Interconnection*, i.e. network expansion, which means having enough transmission capability to bring power from distant sources to satisfy local excess demand.

3. *Reserve capacity of dispatchable sources*, which means holding sources of power such as gas-fired electricity in reserve, that can be quickly deployed when needed.

4. *Optimizing the wind/solar ratio*, which means choosing a mix that best utilizes natural correlations between supply and demand.

5. *Overcapacity*, which means installing capacity so that the annual generation is higher than the annual demand.

6. *Energy storage technologies*, which means storing energy so that it can be used when needed; common methods of energy storage include batteries, pumped-hydro, capacitors, compressed-air, flywheels, gravity-based systems and storing chemical energy in fuels.

We now discuss each of these in more detail.

We do not model demand-side management technologies explicitly because very little data exists on costs and its progress is hard to model (i.e. there is no good cost savings data that allows us to apply Wright's law (see Section 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 believed to be a cheap way of matching supply and demand (Brouwer et al. 2016). 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 provides a good solution when unproductive weather is localized and power can be shipped from locations with productive weather to locations with unproductive weather. Studies often find it to be a cheap solution for the intermittency of renewables, e.g. Brouwer et al. (2016). 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 assume that the level of interconnection will be roughly equivalent in all scenarios, so this cost can be neglected when considering the relative costs of scenarios. Note that the No Transition scenario assumes a doubling of the global electricity supply system by 2050 (as most models do), so there will be a great deal of investment in power grids taking place in any case. Most of the growth will be in developing countries, but existing grids will also require a lot of maintenance and upgrading. We take this to be our interconnection baseline, roughly constant across all scenarios. While our model is designed for simplicity, it is clear that 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. (2018), Denholm & Mai (2019), Solomon et al. (2014), Saarinen et al. (2015), Ziegler et al. (2019)).

Current solutions to matching fluctuating supply and demand in electricity markets depend strongly on maintaining a reserve capacity of dispatchable sources. Since all of our scenarios begin with the present mix, they all make use of this at the outset. However, our rapid renewable 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 fuels that are produced by electrolysis 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 cycle (Shaner et al. 2018). 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 variable renewable energy in the generation mix. In addition, long periods of reduced variable renewable energy 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 can be considered as 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 costs (Blanco & Faaij 2018), 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 requirement, 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 Shaner et al. (2018) 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 as a function of time. In fact, it should be noted that our modeling 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 sources (Lovins 2017).) 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 short term imbalances in the supply and demand of electricity and to provide energy storage for electric vehicles. We assume that Li-ion batteries are used, though 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

$$q_t^{11} = q_t^{11,grid} + q_t^{11,transport}. \tag{14}$$

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 daily variable renewable energy 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 however that β_{ST} is the fraction of electricity *generated* by wind and solar, not the fraction of end-use electricity *delivered* by wind and solar via the grid. In a high renewables scenario variable renewable energy generation will be much larger than variable renewable energy delivered via the grid, because so

much electricity is required in the intermediate energy sector for P2X fuel production (perhaps 50% larger, see Panel H of Figure S7).

To illustrate the point, suppose that total generation is exactly 50% larger 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, so leaves plenty of room for maneuver.

Solomon et al. (2014) 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. Variable renewable energy 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 this 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

$$q_{t+1}^{11,grid} = \min\left\{(1 + g_{ST}^{grid})q_t^{11,grid},\ \frac{\beta_{ST}}{365}(q_{t+1}^{9} + q_{t+1}^{10})\right\}. \tag{15}$$

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. On the one hand, the ability to charge at multiple locations throughout the day reduces the size of battery pack required for a given daily usage, but on the other, Li-ion batteries currently perform best when operated within charge a range of around 25-85%, so a larger battery pack is required for a given daily usage.

However, given our one-full-cycle-per-day assumption, this implies that there must be $\frac{1}{365}EF_t^{4,1}$ EJ of EV battery pack capacity on the system. 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

$$q_{t+1}^{11,transport} = \min\left\{(1 + g_{ST}^{transport})q_t^{11,transport},\ \frac{1}{365}EF_{t+1}^{4,1})\right\}, \tag{16}$$

and q_t^{11} is now fully determined.

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 good economies of scale[9] (see e.g. Ha & Gallagher (2015), Li et al. (2017)). 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, and 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

[9]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.

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

$$q_{t+1}^{12} = \min\left\{(1 + g_{LT}^{grid})q_t^{12},\ \frac{\beta_{LT}}{365}(q_{t+1}^{9} + q_{t+1}^{10})\right\}. \tag{17}$$

In the example above (Section 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. This is not a large amount of storage, but we include it to increase the diversity of options available.

3.3 P2X fuels for long term energy storage

We use *Power-to-X* (P2X) fuels as a representative technology for long term energy storage. 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. 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. One is for end use, to provide a substitute for fuels such as diesel fuel in transportation, where portable fuels are needed, or coal in applications where heat is needed, such as steel manufacture. The other application is to produce electricity during unproductive weather.

There are currently two leading technologies available for converting hydrogen to electricity: fuel cells and direct combustion in gas turbines. 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.[10] Ammonia can also be combusted in turbines, though combustion must be complete in order to avoid undesirable combustion products (Kobayashi et al. 2019).

However, for model simplicity we suppose that reversible electrolyzer/fuel-cells are used, 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 options currently available (AEL, PEM, SOEC — see Section 6.12), 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 is expected to 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, and 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 will be much larger than the quantity for grid backup, so the requirement for fuel-cell capacity will not be a limiting a factor.[11]

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 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 IEA (2019*b*), WEC (2018), Blanco & Faaij (2018).

[10] All EU turbine manufactures have committed to 100% hydrogen capability by 2030 (EUTurbines 2019).

[11] Alternatively we could assume that gas combustion turbines are used for generating electricity from hydrogen. This would require adding an extra technology to represent the capital and operating costs of gas turbines (*not* the hydrogen fuel costs though, as these are already accounted for), but we prefer to avoid this extra complication. However, there is a currently a very large fleet of gas turbines, which will likely be required for at least the next 15 years in all scenarios, so if after this time gas-fired electricity is phased out there will be plenty of capacity available to burn hydrogen, at very low cost. Most likely a combination of reversible fuel cells and combustion turbines would 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.

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-hydro storage). 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 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 or methane may be a good solution.

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 $/kWh (Ikäheimo et al. 2018), 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 zero carbon industrial hub.

The total cost of transmission of hydrogen by pipeline is around double that of natural gas (Icerman 1974), per unit of energy. 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 gas (Brown et al. 2011) and hydrogen (Ogden 2018), as they depend mostly on labour, terrain and location. Construction costs vary greatly due to these three components alone (van der Zwaan et al. 2011). Modern natural gas distribution networks can also be used to transport hydrogen with only minor upgrades (Dodds et al. 2015, Staffell et al. 2019). 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 then the cost of new pipelines (Speirs et al. 2018). New ammonia pipelines are cheaper than new hydrogen pipelines.

For longer (intercontinental) distances, ammonia is already transported by sea in chemical 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 (presumably at similar, relatively high cost), and the first such ship was launched in Japan in 2019 (Harding 2019).

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. (2015) and Butera et al. (2019) 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 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.

The P2X sector in our Fast Transition scenario is, in 2070, over 2.5 times smaller than the fossil fuel sector, so it is safe to assume that infrastructure installation, modification and upgrade costs for P2X fuels are at most similar to the maintenance, upgrade and expansion costs that would be incurred anyway in a 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 in the electricity mix in the long run for wind and solar 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 end-use 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 end-use electricity is X EJ, then a maximum of $(\psi^9 + \psi^{10}) \times X$ EJ of end-use electricity comes from wind plus solar. We assume that $(1-\theta) \times (\psi^9 + \psi^{10}) \times X$ EJ comes directly from generation, but the remaining $\theta \times (\psi^9 + \psi^{10}) \times X$ EJ must be generated from P2X fuels *made from these sources*. We must 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 end-use 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 face (Cannon et al. 2015). Grams et al. (2017) 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. 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, but nevertheless assume that multi-week lulls in variable renewable energy output may occur, so specify that in a very high renewables scenario there must be enough P2X fuel stored to cover one month's worth of end-use 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 end-use 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

$$\theta\left(\psi^9 + \psi^{10}\right) EF_t^{4,total} \min\left\{\Pi_t^{P2X\,grid}, 1\right\}. \tag{18}$$

Both methods currently have a conversion efficiency of around 50% (combusting hydrogen in gas turbines has the same efficiency as when burning natural gas, and see Staffell et al. (2019) for fuel cell efficiencies), so the quantity of P2X fuels required is twice the quantity of electricity given in Eq. (18). 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:

$$P2X_t = \underbrace{EF_t^{5,total}}_{\text{Direct use}} + \underbrace{2 \times \theta\left(\psi^9 + \psi^{10}\right) EF_t^{4,total} \min\left\{\Pi_t^{P2X\,grid}, 1\right\}}_{\text{Electricity backup}}. \tag{19}$$

3.4 Verifying electricity system reliability

Shaner et al. (2018) performed a detailed study assessing the geophysical limits on solar and wind generation in the US, based on 36 years of weather data. They investigated how grid reliability depends on (1) the wind/solar generation mix, (2) annual wind/solar overgeneration (above annual final electricity demand), (3) storage of varying durations, and (4) interconnection (aggregation of resources over large areas). One important feature of the work was that it assumed perfect transmission and storage of energy, with no losses or constraints — the results are therefore not as directly applicable as desired. In particular, including energy losses would mean that more installed generation capacity is required to meet energy demand in the periods of lowest resource availability. This would in turn cause higher generation at all other times, increasing the annual overgeneration, counterintuitively *reducing* storage and interconnection requirements. Their results may therefore overestimate storage and interconnection requirements for a given level of overgeneration. Nevertheless, the results likely provide a 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. We assume that the results are applicable globally, i.e. that similar resource variability exists and can be utilised over large enough areas to smooth out renewable energy supply.

The study showed that solar and wind can quite easily meet up to 80% of final electricity demand, with only minimal levels of excess energy generation or interconnection or storage. In this case the remaining 20% of demand must be met by dispatchable generation sources. However, providing *either* enough storage capacity to store 50% of the average daily final electricity demand *or* enough solar and wind generating capacity such that their annual electricity generation is 150% of the aggregate annual final electricity is sufficient to reliably serve 85% of final electricity demand (under existing levels of interconnection). Combining both of these capabilities results in the ability to meet around 97% of final electricity demand reliably.

The most challenging situation we consider is the Fast Transition scenario (Figure S7, Panels E-J). From 2050 onwards in this scenario there is enough dedicated daily grid storage capacity to store 40% of daily final electricity (daily grid batteries plus multi-day storage — the latter can of course be cycled daily as required.). However, including the stock of EV batteries takes the total daily storage to around 60% of daily final electricity. In addition solar and wind generation reaches 100% of final electricity demand by 2040, 125% in 2050, then 139% after this.[12] Thus enough solar, wind and battery storage technology is present to achieve grid reliability at around the 90-99% level. However, variable renewable energy sources in the Fast Transition scenario only account for 82% of the generation mix in 2040 (then 85% in 2050 and up to 89% after that) so there are many dispatchable generation sources available to ensure reliability well above current regulated requirements (99.97% in the US).

In addition to this already sufficient level of technology deployment we assume that enough P2X fuel is available to cover one month's worth of solar and wind's combined contribution to final electricity. Performance requirements for bulk energy storage in grid applications are significantly lower than for transport, but in our model we assume a 10 year lifespan for all daily batteries. They will continue to function at lower performance for much longer than this though, perhaps another 10 years, so we anticipate very large quantities of second-hand batteries also being available for grid applications. Thus both our cost and reliability estimates are conservative.

For a different perspective we now compare this to the evidence presented in Blanco & Faaij (2018). They review 35 different studies on energy storage requirements. There is a wide range of different views. To reach a variable renewable energy penetration level of 80-94%, the estimates in these studies range from 0.01% to 1.6% of annual final electricity. (The percentages refer to the fraction of annual energy use by electricity that must be stored). For variable renewable energy penetrations higher than this, 80% of studies in the literature suggest that storage ranging from 0-4% of annual final electricity is sufficient (less than 15 days per year of final electricity), rising to 7% (26 days) if 90% of studies are considered. By comparison, we are assuming one month of P2X storage, corresponding to 8.3%; only three of the thirty five studies estimate that more than this is required. However, we are also assuming substantial daily storage and three day flow battery storage; if these are included then our planned storage surpasses even the most pessimistic estimates.

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 satisfies the energy generation 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, overgeneration 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.

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. (2011) and Rasmussen et al. (2012) 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. (2014) 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

[12] 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.

and energy demands.[13]

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 (IEA 2019*b*), and the expected future values are around 70% and 90% respectively. 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 required by the system is the sum of electricity used by sectors and that used for electrolysis:

$$Elec_t = \underbrace{EF_t^{4,\,total}}_{\text{Direct use}} + \underbrace{1.43 \times P2X_t}_{\text{Electricity backup}} . \tag{20}$$

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 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% capacity factor, so they operate for half of all hours each year.[14] Therefore, in order to produce the quantity $P2X_t$ EJ of P2X in year t (given in Expression 19), the total capacity of electrolyzers available on the system must be $2 \times P2X_t$ EJ/year. Electrolyzers are assigned technology index 13 in the model, so, in units of EJ/hour we have

$$q_t^{13} = \frac{2}{24 \times 365} P2X_t \tag{21}$$

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.

Our electrolysis capacity factor assumption is probably reasonable (if not even pessimistic) because there will likely be more than enough battery storage on the system to allow load-shifting over hours and days to enable the electrolyzers to operate for at least 50% of the time on average.

3.7 Technology lifespan and annual additions

In order to calculate the costs of many of the technologies 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, due to the way we define technologies, "quantities" are in units of energy (to maintain the simple 'energy-only' approach), but for most technologies we use annuitized (levelized) costs, which persist for the entire lifespan of the underlying capital stock. So 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 $q_t^i = \max\left\{0, q_1^i(1-(t-1)/L^i)\right\}$. However, accounting

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

[14] We pick the central value here for parsimony. However, note that this electrolysis is powered by electricity generation that is added to the model specifically for this purpose, so it is likely that electrolyzers would be placed in locations with the cheapest, most reliable combinations of wind and solar, which may allow a higher average capacity factor.

for new additions on top of this initial decaying production gives

$$q_t^i = \sum_{\tau=1}^{t} Q_{t,\tau}^i \tag{22}$$

$$Q_{t,\tau}^i = \begin{cases} \max\left\{0,\ q_1^i\left(1-\frac{t-1}{L^i}\right)\right\} & \text{for } \tau = 1 \\ \max\left\{0,\ q_t^i - \sum_{s=1}^{t-1} Q_{t,s}^i\right\} & \text{for } \tau = t \text{ and } t > 1 \\ Q_{t-1,\tau}^i & \text{for } 1 < \tau < t \text{ and } 1 < t \leq L^i \\ \max\left\{0,\ Q_{t-1,\tau}^i - Q_{t-L^i,\tau}^i\right\} & \text{for } 1 < \tau < t \text{ and } t > L^i \end{cases} \tag{23}$$

Eq. 22 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. 23, 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 S4 shows how the process works.

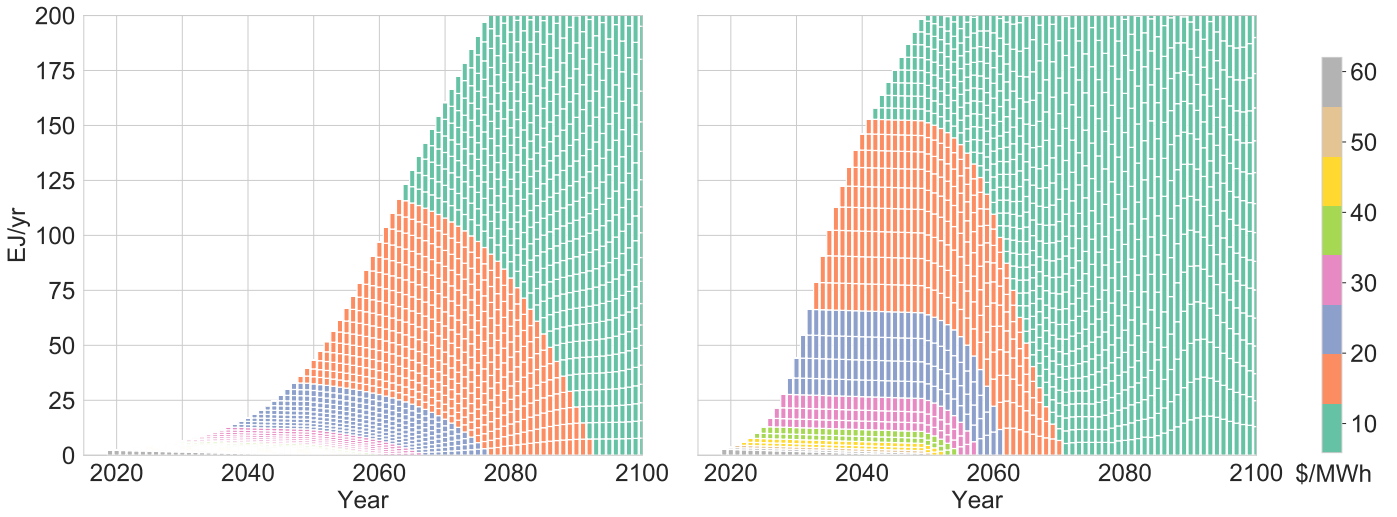

Figure S4: Annual solar PV generation broken down by PV system vintage in a slow deployment scenario, left (initial annual growth 10%), and a fast deployment scenario, right (initial annual growth 30%). Colours show the median forecast LCOE in the deployment year. For any given year, the block of energy 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.

3.8 Scenario construction summary

We have thus determined q_t^i for all i and t. Following the above steps a scenario is created by specifying *i*) the final carrier mix in each sector, and the carrier growth rates for getting to there from the current mix, *ii*) the final electricity mix and the technology growth rates for getting to there from the current mix, and *iii*) the initial growth rates of batteries for ramping up storage capacity to levels required for large quantities of wind and solar. Scenarios may be generated fairly smoothly in this way, but it is still a manual, ad-hoc process that requires care to avoid unrealistically sharp transitions or fluctuations in slack variables.

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 & Haigh (2009) 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". First, this does not explain the 1% threshold, but more importantly, it says nothing about how growth proceeds when total costs *do* drop below existing technology operating costs, as we predict they will soon. 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 the current rate should occur in roughly 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 four scenarios, listed below. In all cases useful energy in each sector, and economy-wide, grows at 2% until 2100. All technology growth rates are less than or equal to their current (or very recent) observed rates, apart from two minor exceptions.[15] 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 4.1-4.4 show the scenarios in more detail.

- ***No Transition:*** This scenario represents minimal change to the current energy system. We assume that fossil fuel growth remains at current levels and there is relatively little deployment of renewable energy technologies beyond what already exists. Solar and wind's share of electricity generation grows from 7% to 16% over the next 50 years (this is much slower growth than has occurred historically, but it seems implausible that no growth at all will occur, so we allow for this very modest level). We use this as a baseline to compare to other transition scenarios, as it represents approximately freezing the energy system in its current configuration and assuming pro-rata growth of all technologies to match the 2% increase in total energy demand. See Figures S5 and S6 and Table 4.

- ***Fast Renewable Transition:*** In this scenario, we assume 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 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 electric vehicles.

  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 S7 and S8 and Table 5.

[15] The two exceptions 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 is a transient increase that is just an artifact of the scenario construction process — it could nevertheless be dealt with largely via increased utilization rates of the existing CCGT fleet.

- ***Slow Renewable 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 a lot less P2X fuels backup is required — only one week's cover of solar and wind's contribution is provided. Less P2X fuels is also 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 S9 and S10 and Table 6.
- ***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 Figure S11 and S12 and Table 7.

4.1 No transition

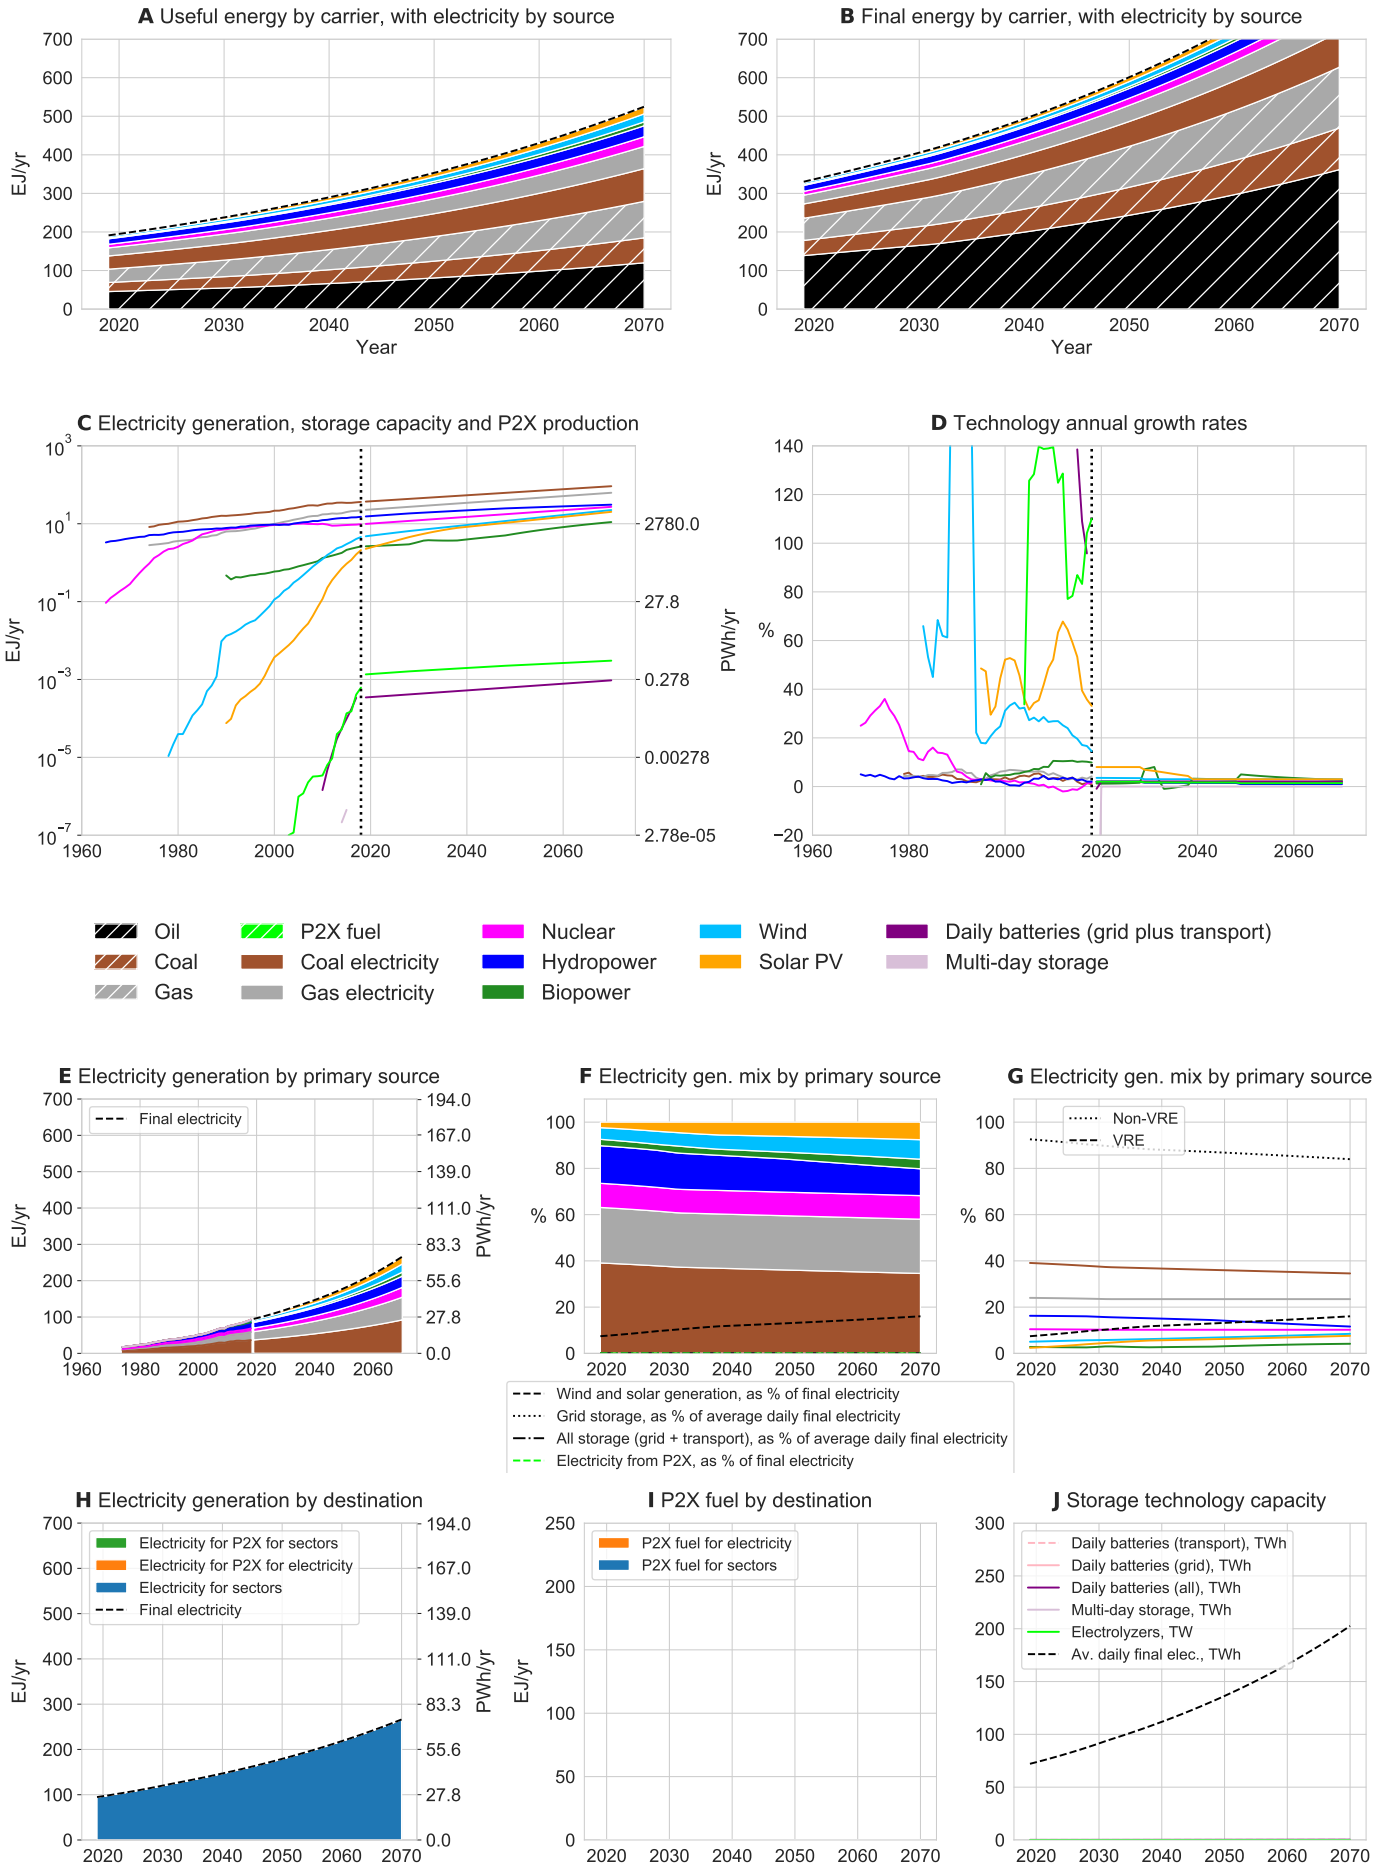

Figure S5: No Transition scenario details. (Historical growth rates shown in panel D are 5-year moving averages, apart from P2X fuel, which is an 8-year moving average.)

Sector	Energy carrier	Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3	Max. % of useful energy
Transport	Oil (slack)	-	-	-	-	-	-
	Electricity	10	10	10	30	1	10
	P2X fuels	2	2	10	30	-	-
Industry	Oil	2	2	10	30	-	-
	Coal	2	2	10	30	-	-
	Gas (slack)	-	-	-	-	-	-
	Electricity	2	2	10	30	-	-
	P2X fuels	2	1	10	30	-	-
Buildings	Oil	2	2	10	30	-	-
	Coal	2	2	10	30	-	-
	Gas (slack)	-	-	-	-	-	-
	Electricity	2	2	10	30	-	-
	P2X fuels	2	1	10	30	-	-
Electricity	Technology						Max. % of generation
Generation	Coal	1.8	1.8	10	30	-	-
	Gas	2	2	10	30	-	-
	Nuclear	2	2	10	30	-	-
	Hydropower	2	1	10	30	-	-
	Biopower (slack)	-	-	-	-	-	-
	Wind	4	3	0	10	1	10
	Solar PV	8	3	10	20	1	10
Storage	Daily batteries	2	0% of daily solar+wind generation stored ($\beta_{ST} = 0$)				
	Multi-day storage		0% of 3-day solar+wind generation stored ($\beta_{LT} = 0$)				
	P2X fuels		No solar+wind grid backup ($\theta = 0$)				

Table S6: No Transition construction parameters.

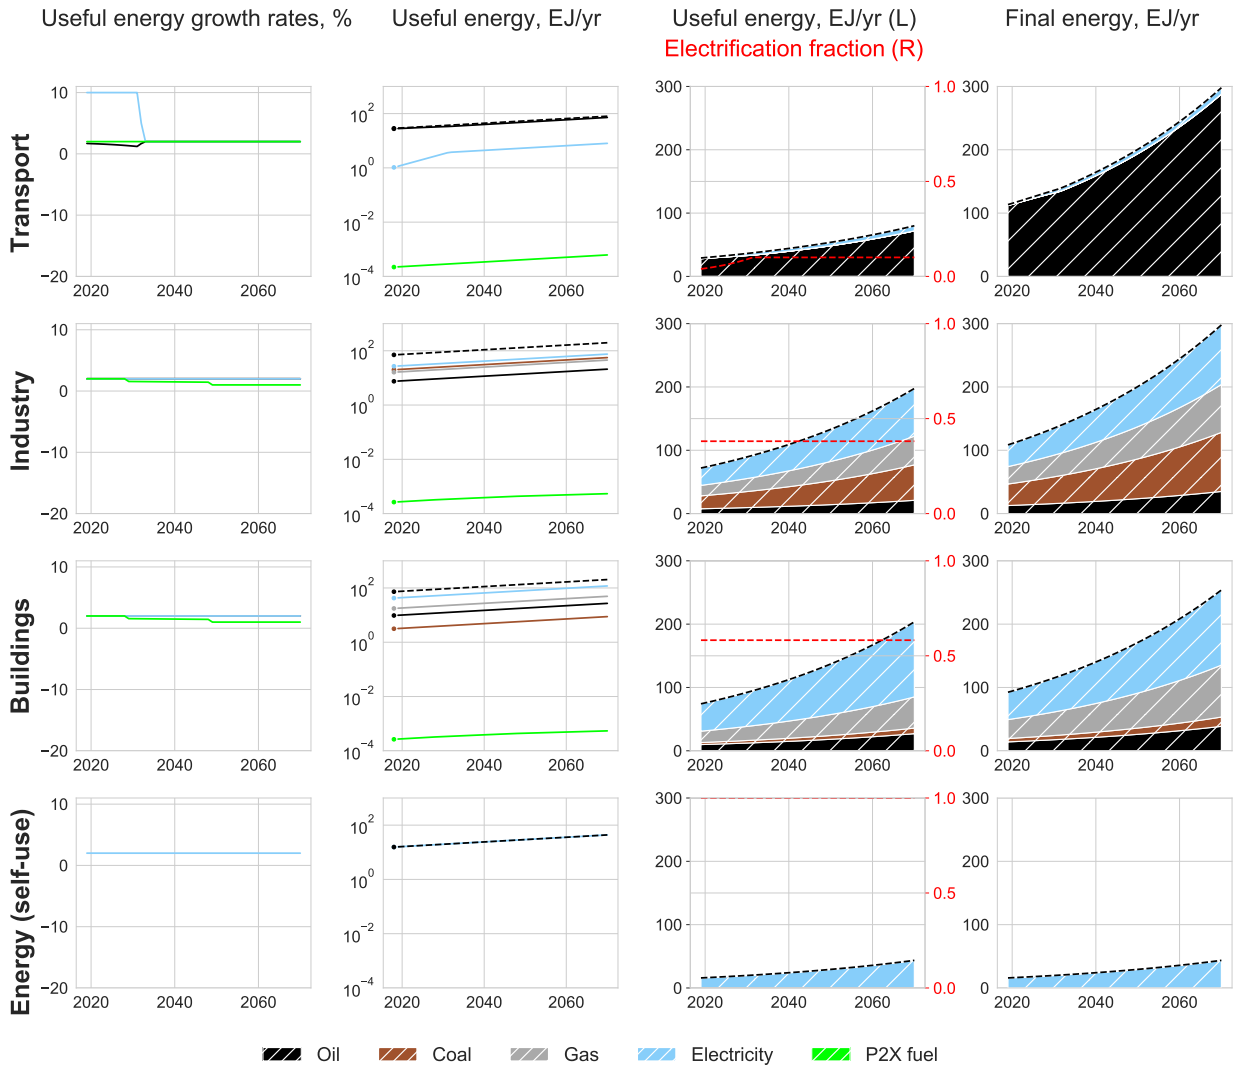

Figure S6: No Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

4.2 Fast transition

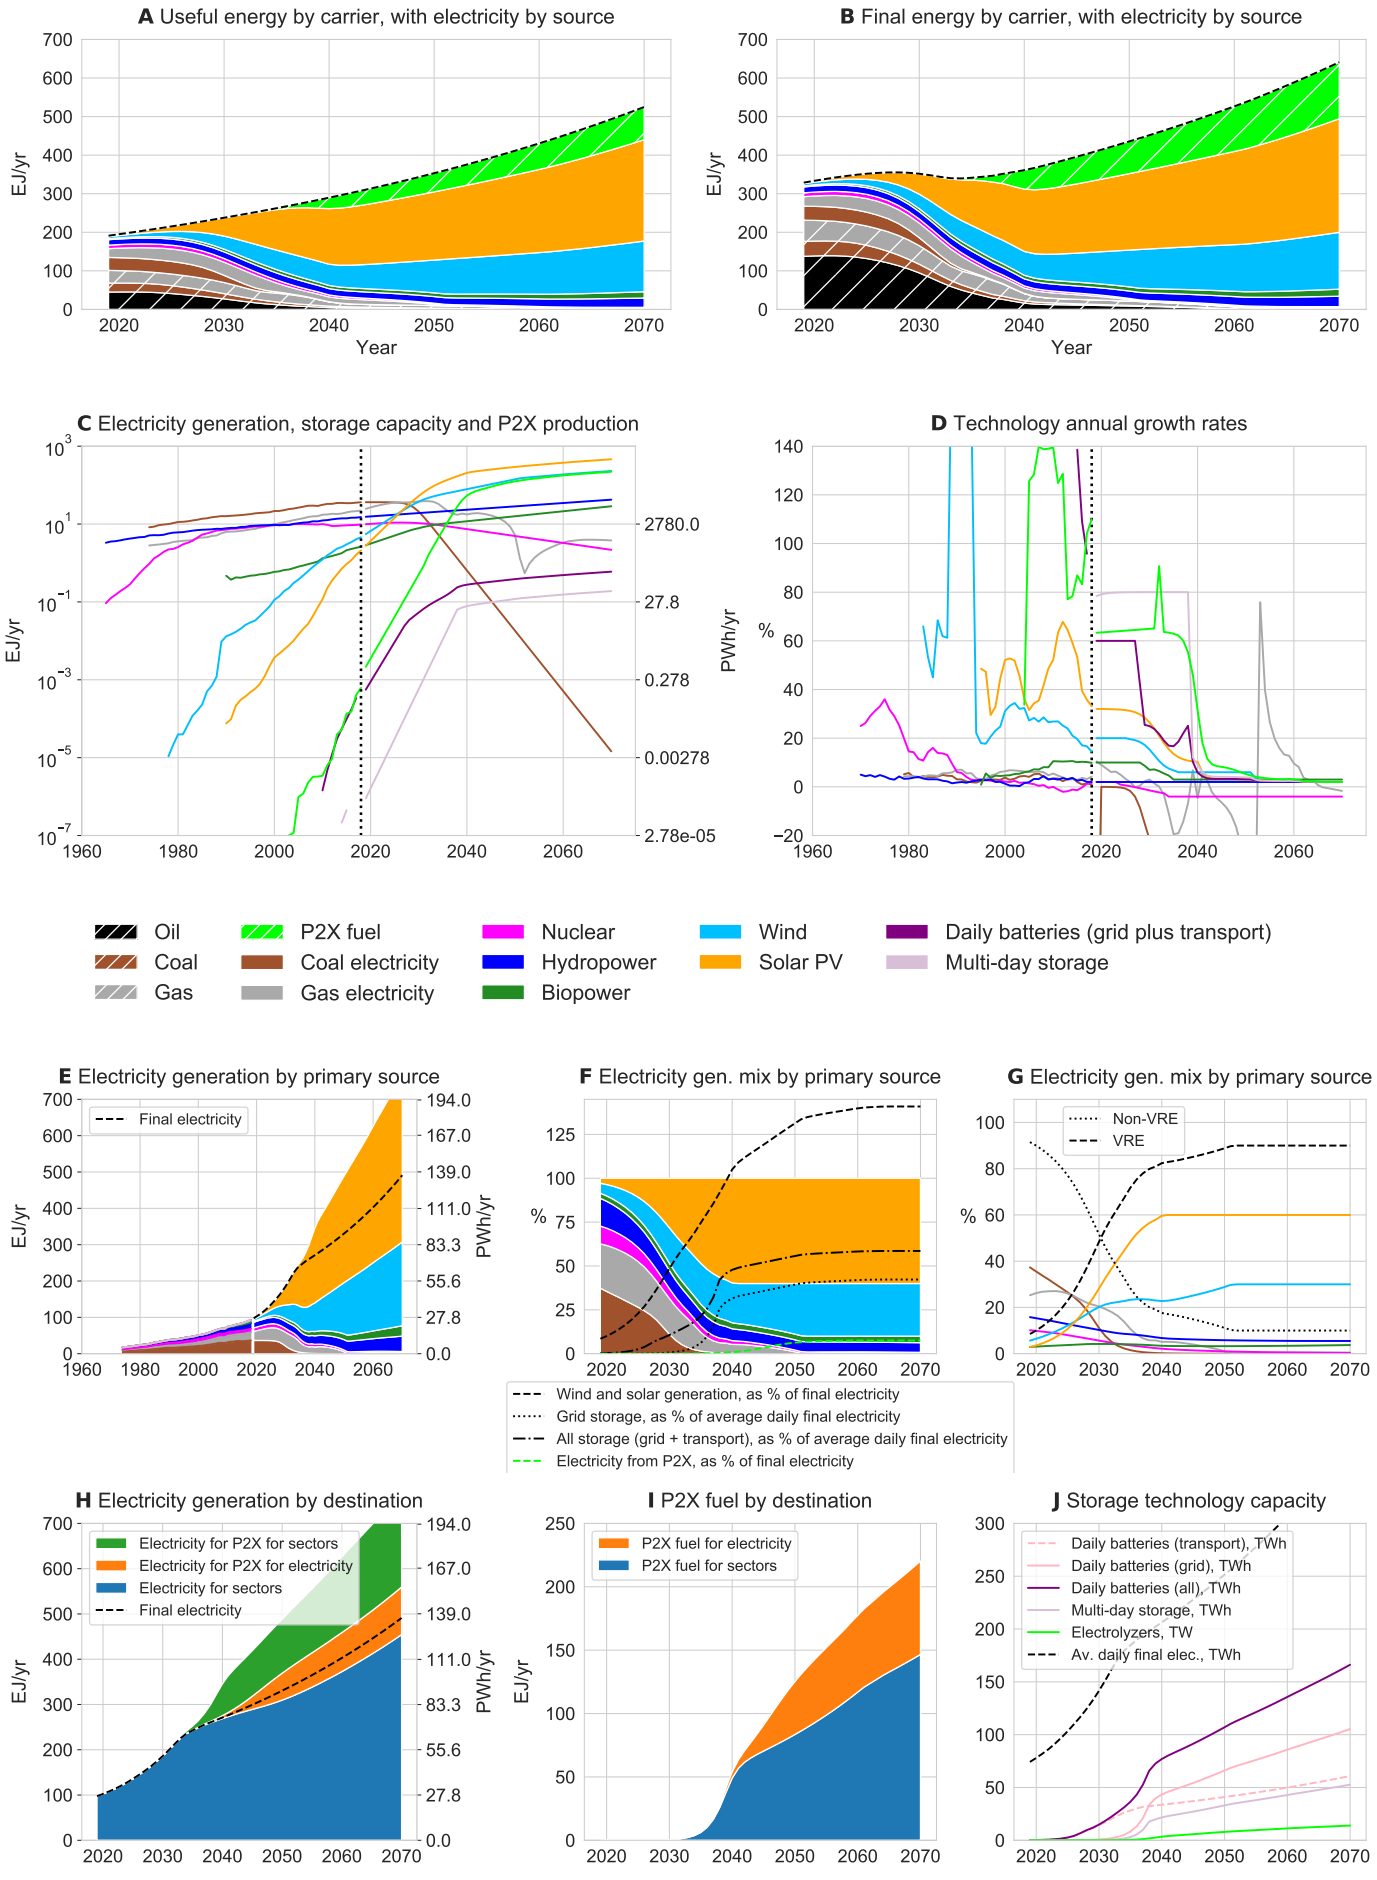

Figure S7: Fast Transition scenario details.

Sector	Energy carrier	Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3	Max. % of useful energy
Transport	Oil (slack)	-	-	-	-	-	-
	Electricity	25	5	8	20	1	80
	P2X fuels	60	4	12	30	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	1	40	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	12	30	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	20	6	7	17	1	30
	Solar PV	32	10	3	23	1	60
Storage	Daily 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 S7: Fast Transition construction parameters.

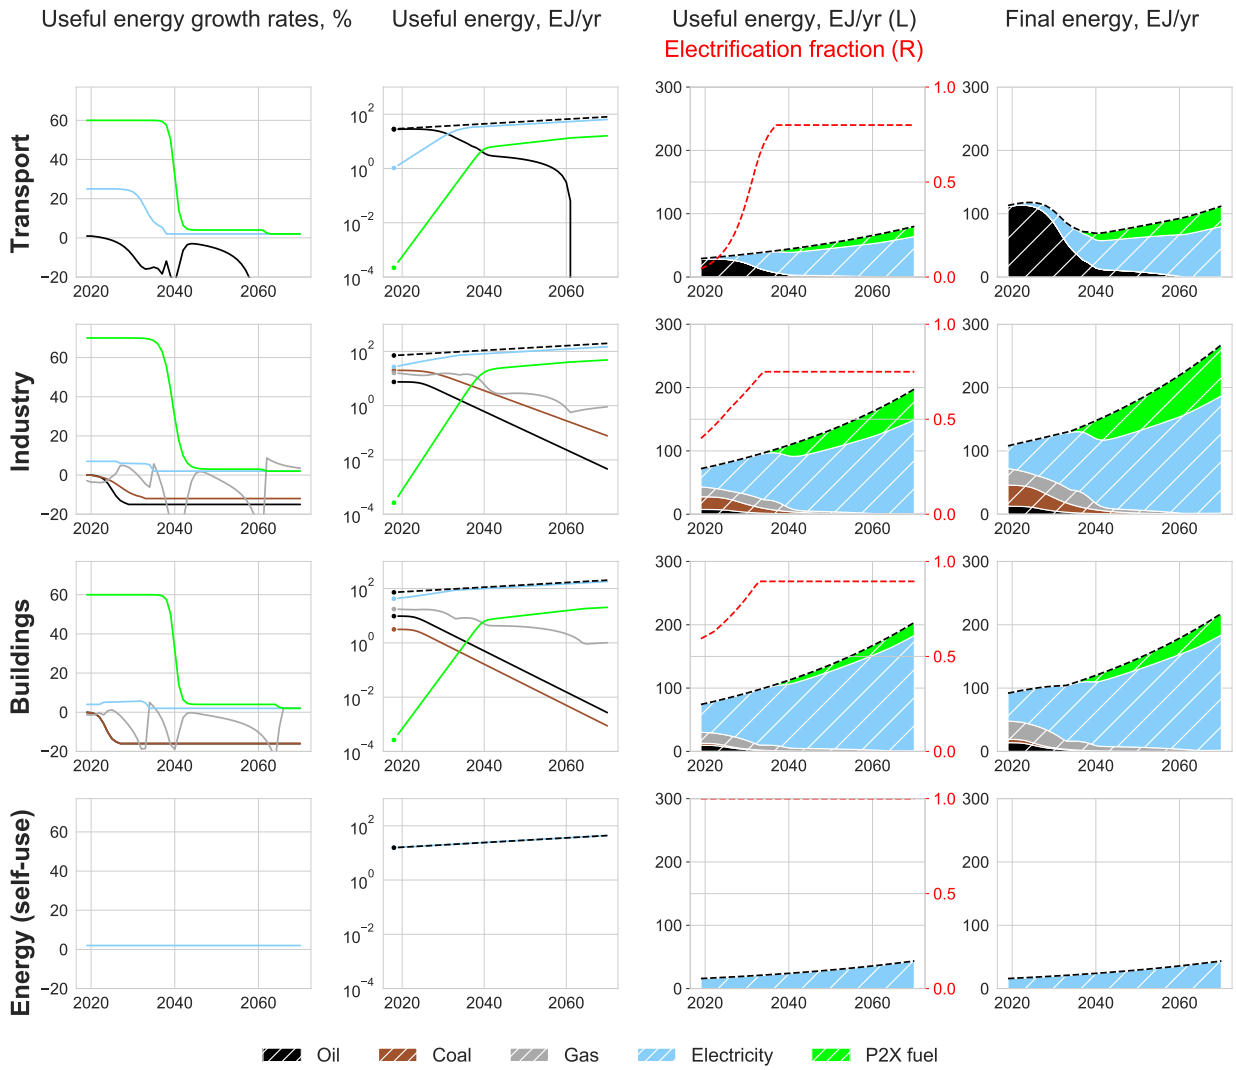

Figure S8: Fast Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

4.3 Slow transition

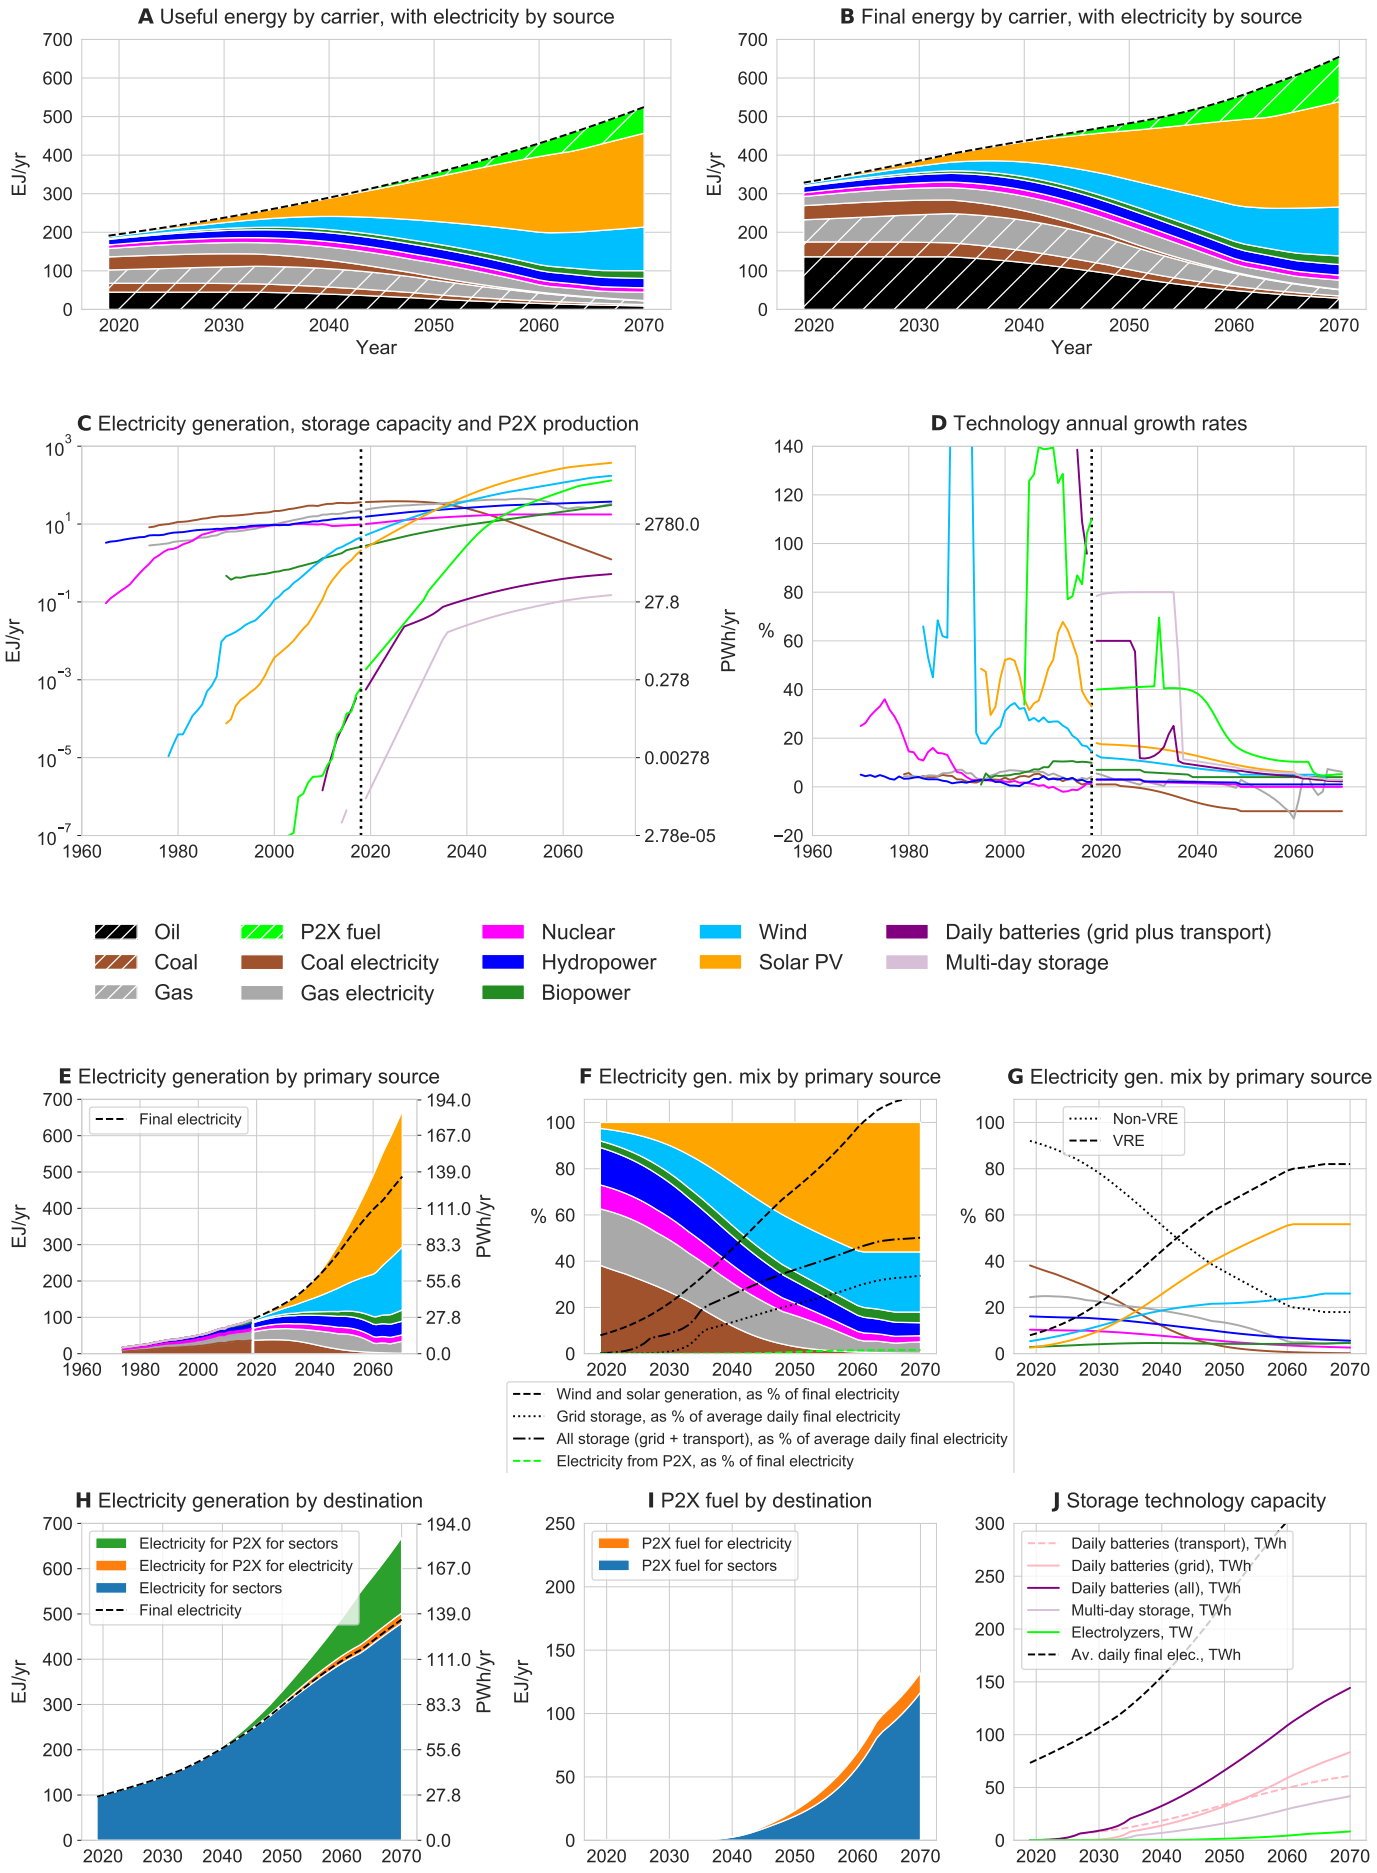

Figure S9: Slow Transition scenario details.

Sector	Energy carrier	Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3	Max. % of useful energy
Transport	Oil	0	-5	15	30	-	-
	Electricity (slack)	-	-	-	-	-	-
	P2X fuels	35	12	10	40	1	18
Industry	Oil	0	-6	15	30		
	Coal	0	-6	10	30	-	-
	Gas	2	-6	15	40	-	-
	Electricity (slack)	-	-	-	-	-	-
	P2X fuels	45	12	10	40	1	20
Buildings	Oil	0	-6	15	30		
	Coal	0	-5	15	30		
	Gas	2	-9	15	40	-	-
	Electricity (slack)	-	-	-	-	-	-
	P2X fuels	40	10	10	40	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	13	5	1	30	1	26
	Solar PV	18	5	1	46	1	56
Storage	Daily 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 S8: Slow Transition construction parameters.

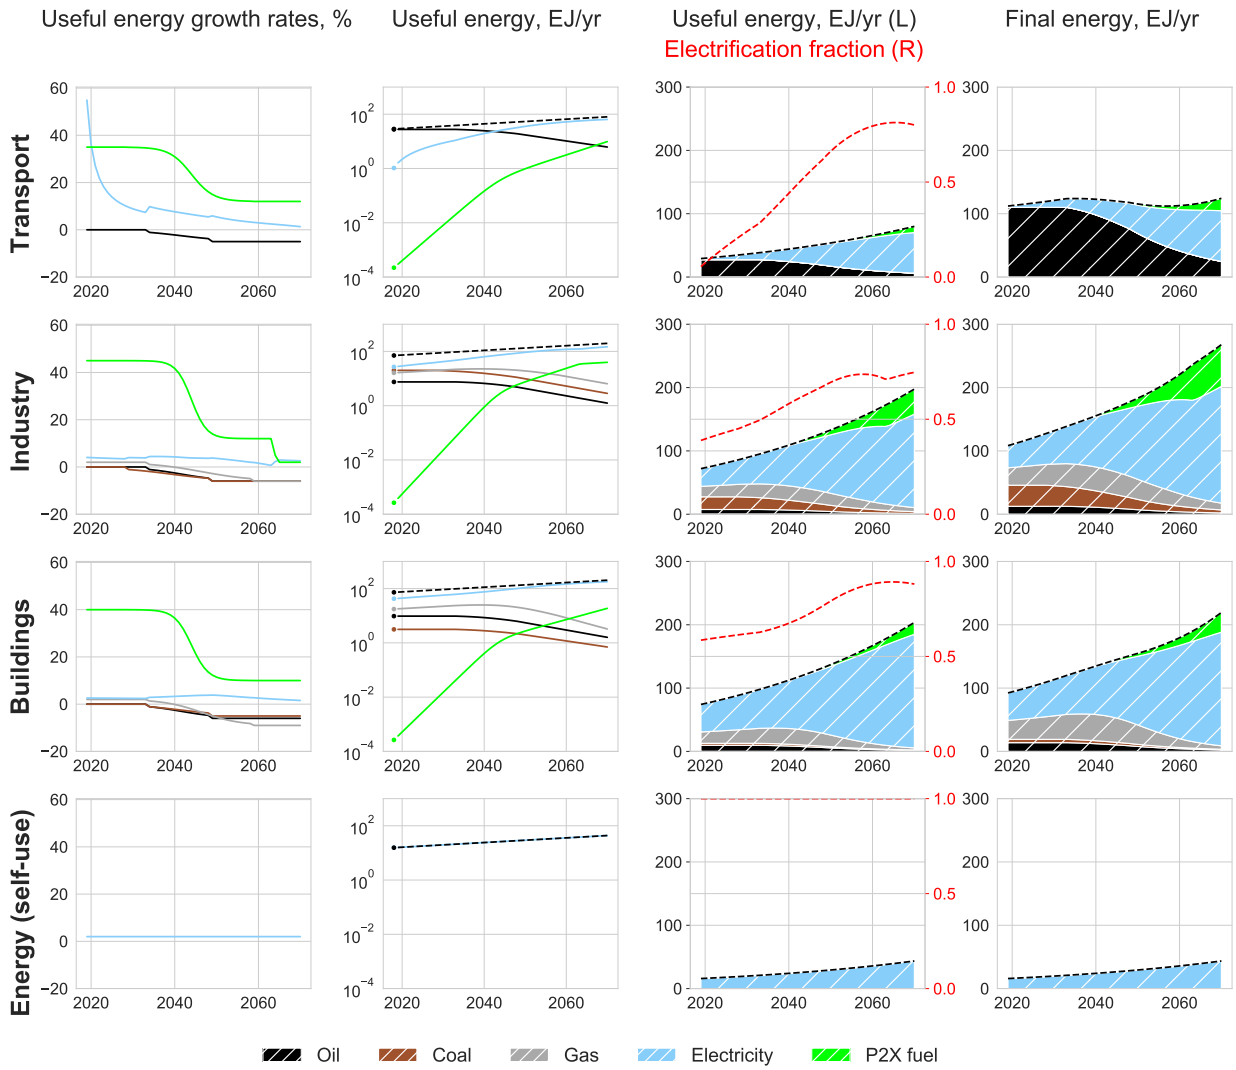

Figure S10: Slow Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

4.4 Slow nuclear transition

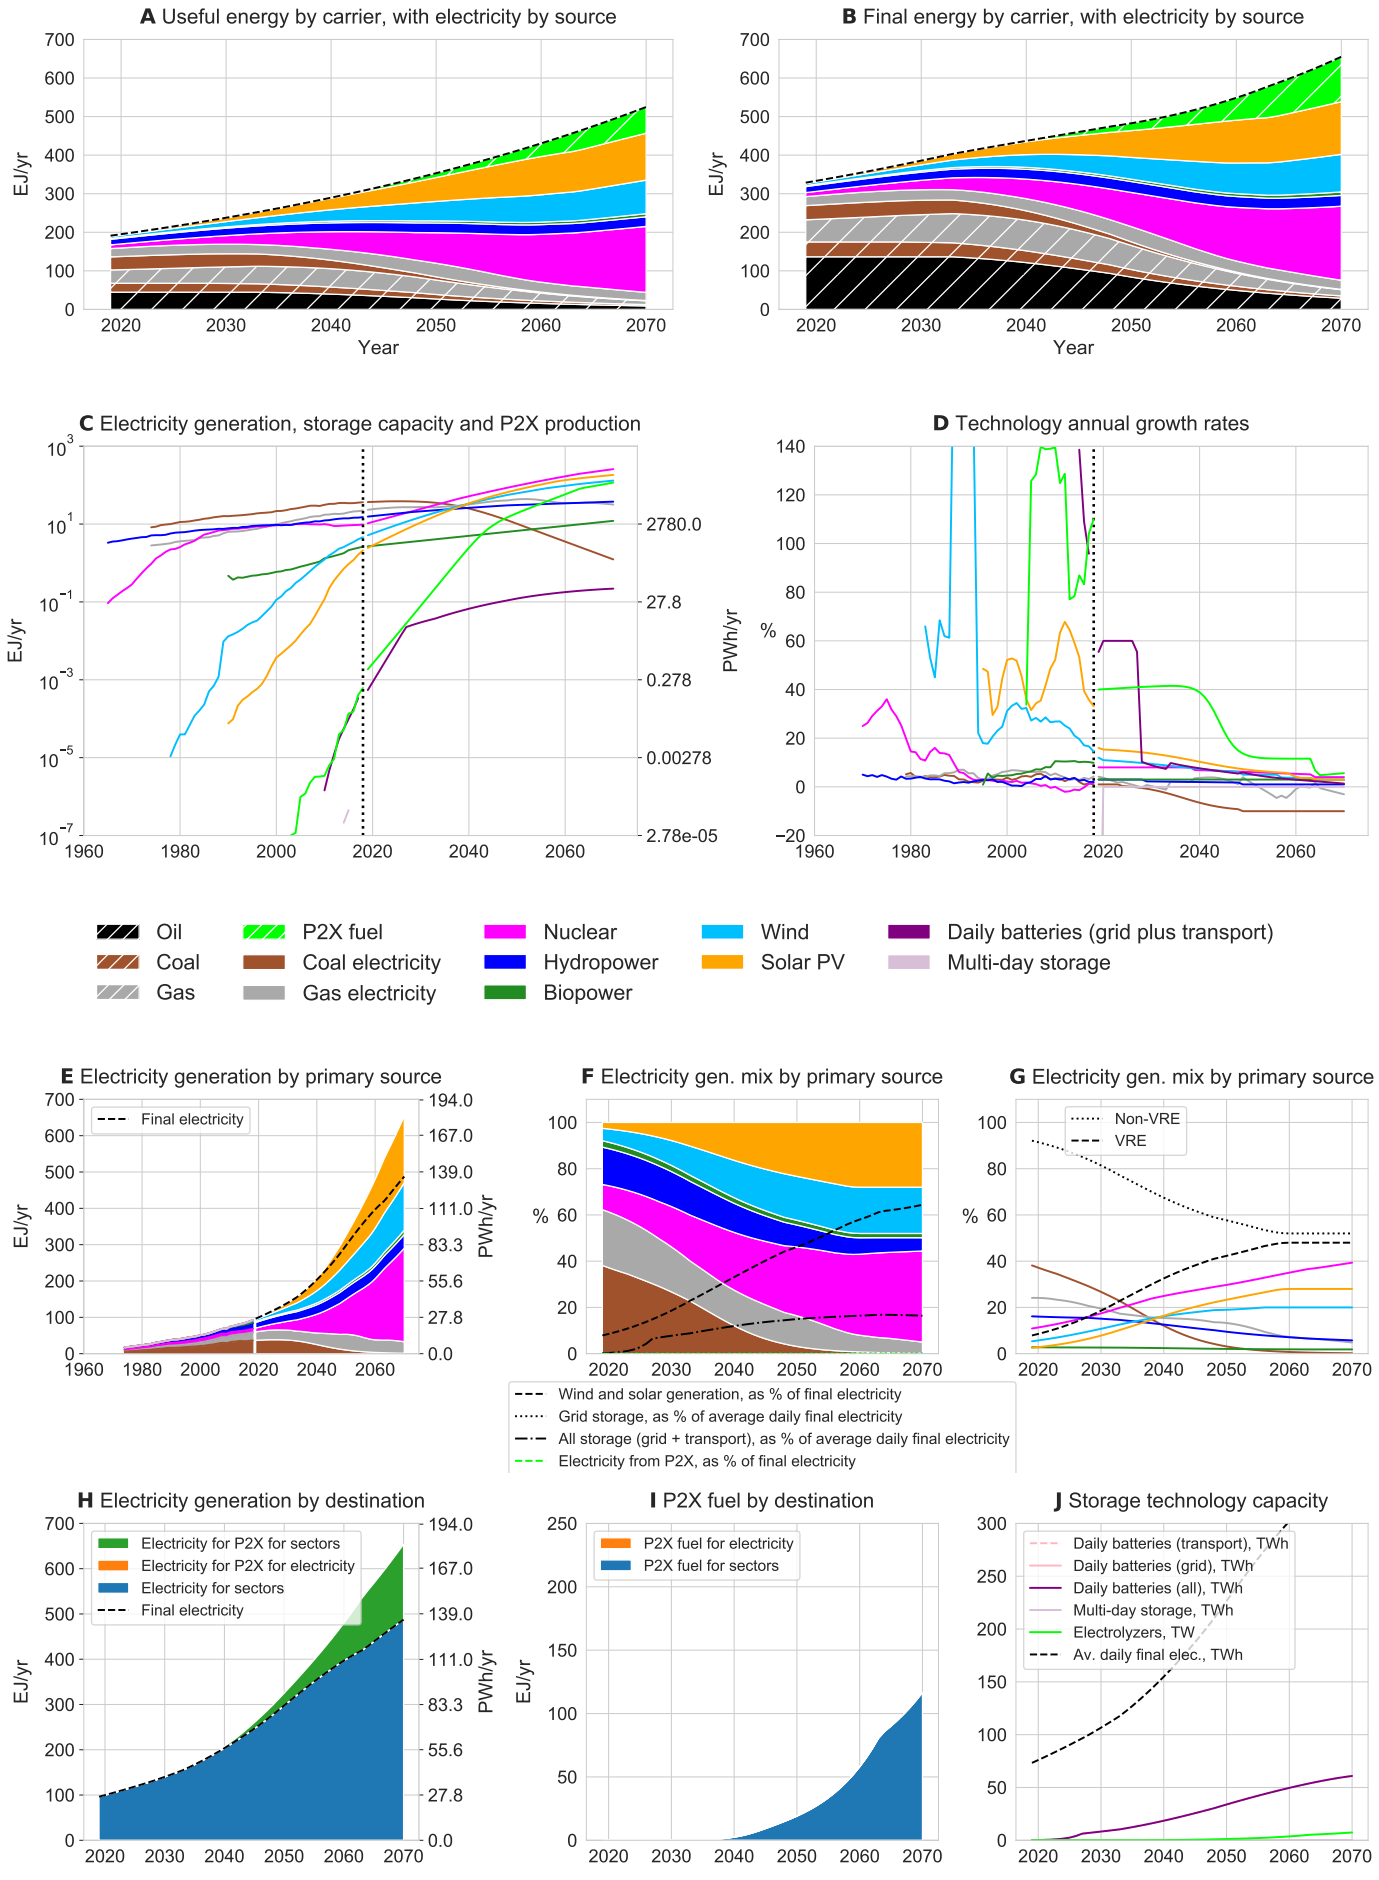

Figure S11: Slow Nuclear Transition scenario details.

Sector	Energy carrier	Initial growth rate, %	Final growth rate, %	t_1	t_2	t_3	Max. % of useful energy
Transport	Oil	0	-5	15	30	-	-
	Electricity (slack)	-	-	-	-	-	-
	P2X fuels	35	12	10	40	1	18
Industry	Oil	0	-6	15	30		
	Coal	0	-6	10	30	-	-
	Gas	2	-6	15	40	-	-
	Electricity (slack)	-	-	-	-	-	-
	P2X fuels	45	12	10	40	1	20
Buildings	Oil	0	-6	15	30		
	Coal	0	-5	15	30		
	Gas	2	-9	15	40	-	-
	Electricity (slack)	-	-	-	-	-	-
	P2X fuels	40	10	10	40	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	12	5	1	30	1	20
	Solar PV	16	5	1	40	1	28
Storage	Daily batteries	60	0% of daily solar+wind generation stored ($\beta_{ST} = 0$)				
	Multi-day storage		0% of 3-day solar+wind generation stored ($\beta_{LT} = 0$)				
	P2X fuels		No solar+wind grid backup ($\theta = 0$)				

Table S9: Slow Nuclear Transition construction parameters.

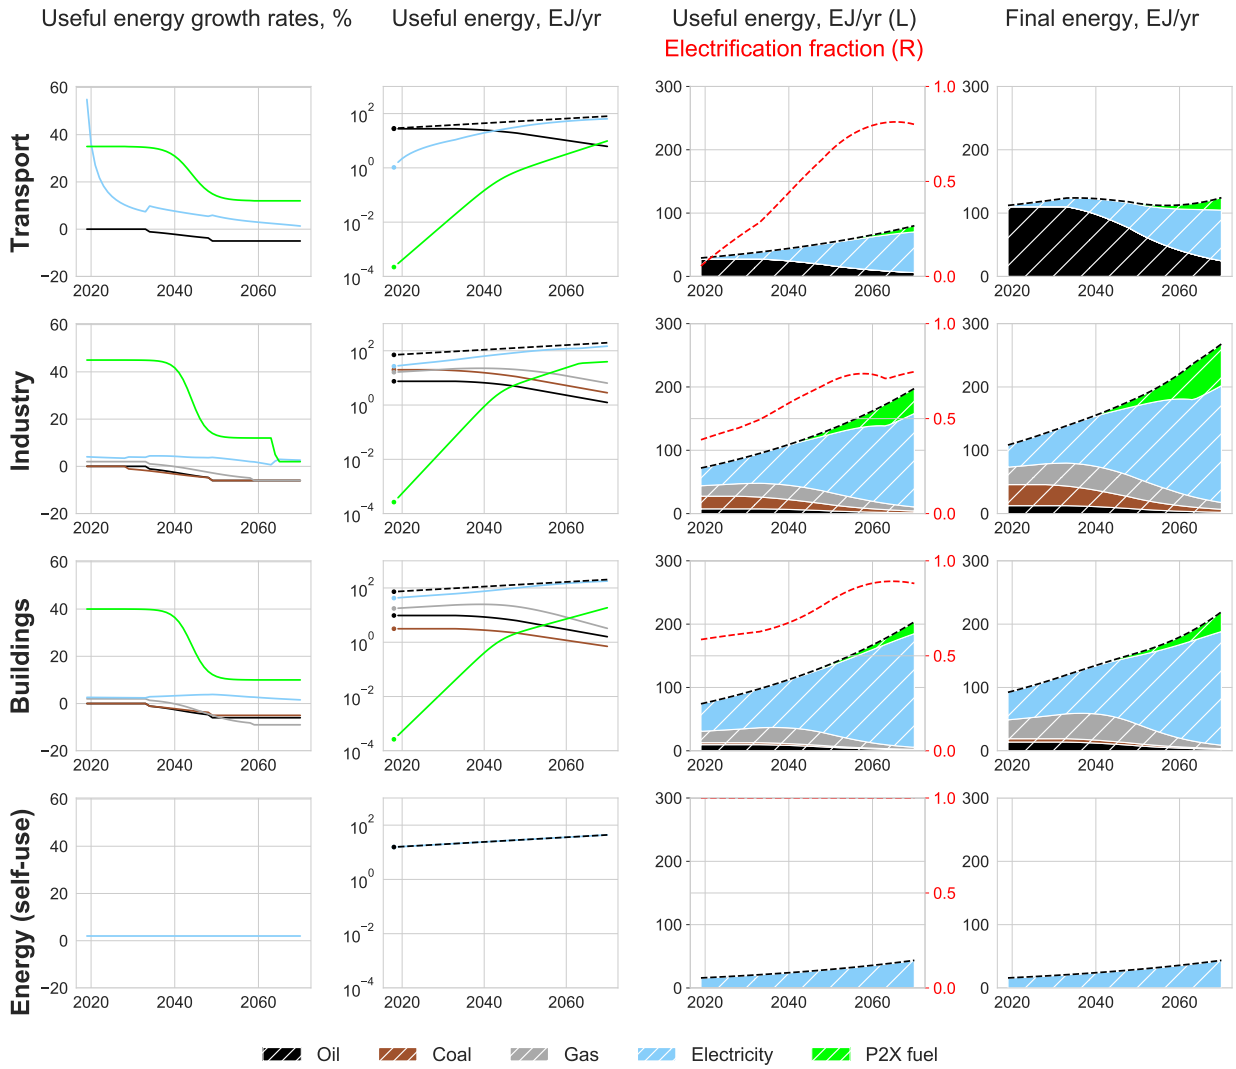

Figure S12: Slow Nuclear Transition carriers by sector. Dashed black lines show exogenous useful energy constraints. Dashed red lines show the extent to which each sector is electrified at each time.

5 Technology cost models

In this section we discuss our calculation of the present discounted cost of each scenario. 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 Section 7.4 we show the effects of using Moore's law instead of Wright's law.)

An extensive literature exists on the experience curve concept and its place within the technological change and innovation literatures. For an overview of its origins, theory and limitations see e.g. Thompson (2012), Arrow (1962), Mazzola & McCardle (1997), Nordhaus (2014) and Way et al. (2019). For details of applications in the field of energy technologies see e.g. Neij (1997), Grubler et al. (1999), Rubin, Azevedo, Jaramillo & Yeh (2015), Samadi (2018), Wene (2016), Ferioli et al. (2009) and Wilson et al. (2020). For a discussion of caveats and subtleties around experience curves see Section 5.6.

5.1 AR(1) process

While being unpredictable in the short-term, real fossil fuel prices have remained remarkably constant — within less than an 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. While there have been dramatic technological improvements in fossil fuel industries at all stages of the production process, 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.

Pindyck (1999) analyzed 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 & Topal (2010) 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 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 S14). An equally plausible model for oil and gas is that they have no long-term trend, but there was a break in the price level in the mid 1970's. We make this choice and calibrate these models from 1975 onward.

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

$$X_t = \phi X_{t-1} + \epsilon_t + \kappa, \qquad \text{with } \phi \in (0,1], \text{ I.I.D. noise } \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2), \text{ and constant } \kappa \neq 0. \tag{24}$$

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,

$$X_t = X_{t-1} + (1-\phi)(\mu - X_{t-1}) + \epsilon_t, \qquad \text{where } \mu = \frac{\kappa}{1-\phi}. \tag{25}$$

$\mu = \mathbb{E}[X_t] = \frac{\kappa}{1-\phi}$ 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 $X_{t-1}, X_{t-2}, \ldots, X_{t-N}$ in Eq. 24. Then

$$X_t = \kappa \sum_{j=0}^{N-1} \phi^j + \phi^N X_{t-N} + \sum_{j=0}^{N-1} \phi^j \epsilon_{t-j}. \tag{26}$$

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

$$X_t = \kappa \sum_{j=0}^{t-1} \phi^j + \phi^t X_0 + \sum_{j=0}^{t-1} \phi^j \epsilon_{t-j}. \tag{27}$$

Thus

$$\mathbb{E}[X_t|X_0] = \kappa \sum_{j=0}^{t-1} \phi^j + \phi^t X_0 \tag{28}$$

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:

$$\mathbb{E}[X_t|X_0] \;=\; \kappa \sum_{j=0}^{t-1} \phi^j + \phi^t X_0 \;=\; \kappa\left(\frac{1-\phi^t}{1-\phi}\right) + \phi^t X_0 \;=\; \mu(1-\phi^t) + \phi^t X_0, \tag{29}$$

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

Using the data shown in Section 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. (2018). 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 randoms 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:

$$\log c_t - \log c_{t-1} = -\omega(\log z_t - \log z_{t-1}) + \eta_t \qquad \text{with IID } \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{30}$$

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 data though (Farmer & Lafond 2016), Lafond et al. (2018) used a version of the model in which the noise shocks are autocorrelated,

$$\log c_t - \log c_{t-1} = -\omega(\log z_t - \log z_{t-1}) + u_t + \rho u_{t-1} \qquad \text{with IID } u_t \sim \mathcal{N}(0, \sigma_u^2), \tag{31}$$

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

$$\sigma_u^2 = \sigma_\eta^2/(1+\rho^2). \tag{32}$$

Lafond et al. (2018) 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. 30) on historical data we obtain estimates for the three values $\hat{\omega}$, $\hat{\sigma}_\omega$ and $\hat{\sigma}_\eta$. Eq. 32 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. 31 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. (2018). 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 $E[c_t] \sim e^{(\hat{\omega}\mu_z + \frac{\hat{\sigma}_u^2}{2})t}$. 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 $Med[c_t] \sim e^{\hat{\omega}\mu_z t}$ 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 $t \to \infty$ if $\hat{\sigma}_u^2 > \hat{\omega}\mu_z$.

Lafond et al. (2018) 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. (2018) 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 a more accurately calibrated model. In practise, most of the series we use have either no or very few data gaps, so the point is mostly theoreticalt.

In Section 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

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

$$C_t^i = c_t^i q_t^i. \tag{33}$$

5.3.2 Wright's law model costs

For each non-fossil fuel based technology ($i = 6, \ldots, 13$) 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 ($i = 6, \ldots, 10$). 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 Section 3.7 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

$$C_t^i = \sum_{\tau=1}^{t} c_\tau^i Q_{t,\tau}^i. \tag{34}$$

This applies for all non-fossil electricity generation technologies.

The three remaining technologies, short-term batteries, mulit-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 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)

$$CRF^i = \frac{r(1+r)^{L^i}}{(1+r)^{L^i}-1}, \tag{35}$$

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, $Q^i_{t,\tau}$ of capacity exists that was installed in year τ, so if the installed system cost in year τ was c^i_τ, then the cost of production of $Q^i_{t,\tau}$ is ($\mathrm{CRF}^i c^i_\tau Q^i_{t,\tau}$) The total cost of producing all technology i, q^i_t, is

$$C^i_t = \sum_{\tau=1}^{t} \mathrm{CRF}^i c^i_\tau Q^i_{t,\tau}. \tag{36}$$

Finally, consider daily-cycling batteries ($i = 11$). These are used in both the power grid and transport sectors (recall Eq. 14: $q^{11}_t = q^{11,grid}_t + q^{11,transport}_t$). The relevant cumulative experience variable in the FDWL model is the total cumulative production of batteries in both applications, q^{11}_t. 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 total quantity of batteries is small anyway, so any inaccuracy in the cost forecast due to omitting electronics production shouldn't affect the results much.

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 — EV batteries are considered just part of routine capital stock turnover. Thus, assuming $q^{11,grid}_t$ and $q^{11,transport}_t$ are broken down by installation year as usual, the cost of production of $q^{11}_t = \sum_{\tau=1}^{t} Q^{11,grid}_{t,\tau} + Q^{11,transport}_{t,\tau}$ is

$$C^{11}_t = \sum_{\tau=1}^{t} \mathrm{CRF}^{11,grid} c^{11}_\tau Q^{11,grid}_{t,\tau} \tag{37}$$

5.4 Total system costs

The total cost of all production at time t in a given scenario is

$$W^{scenario}_t = \sum_{i=1}^{13} C^i_t \tag{38}$$

and the net present cost of the scenario is

$$V^{scenario}(r) = \sum_{t=1}^{T} e^{-rt} W^{scenario}_t, \tag{39}$$

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:

$$NPC^{scenario}(r) = V^{scenario}(r) - V^{No\ Transition}(r) \tag{40}$$

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

$$\mathbb{E}\left[NPC^{scenario}(r)\right] \tag{41}$$

$$\text{and} \quad median\left[NPC^{scenario}(r)\right], \tag{42}$$

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

5.6 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 modeling (for example large errors in solar PV deployment forecasts). Recent work has made progress on understanding causality in learning systems though, and implies that experience is indeed a prominent causal factor in cost reductions in many cases (Lafond et al. 2020). While it is critical to ensure that experience curve patterns of technology development are represented in any model, the limitations of the framework must also be understood. Well known caveats and subtleties include:

- *Technology function, boundaries and aggregation*. Modeling choices must be made regarding how to aggregate 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.

- *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 modeling 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 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 aversion (Way et al. 2019). In practise this method can only be applied to small regions of decision space because nonlinearity 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 four 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 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 support or deployment support 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.

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 Agency (IEA 2019*c*) and BP's Statistical Review of World Energy (BP 2019). Cost data is much harder to find and comes from a wide variety of sources including, among others, Lazard's Levelized Cost of Energy Analyses (Lazard 2019*a*), the International Renewable Energy Agency (IRENA 2019), the U.S. Energy Information Administration's Annual Energy Outlooks (EIA 2019*b*), Bloomberg New Energy Finance (BNEF) and Bloomberg L.P. (via Bloomberg Terminal).

There is no single natural metric for measuring "the cost" of any energy technology (or any technology in general), because each technology has many different properties that together determine how it is used as part of a wider technology system. 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. The LCOE calculation depends on several parameters including capital cost, fuel cost, capacity factor, project lifetime, interest rate etc. (Schlömer et al. 2014). While other studies forecast total energy costs for a technology by forecasting each of these factors separately, due to our focus on global system costs it makes more sense here to make a single LCOE forecast, rather than several separate forecasts, which would then need to be recombined to achieve the same result. Furthermore, since we forecast cost distributions for each technology, calibrating the model on historical data automatically accounts for the historical variation in each separate cost component (e.g. PV module efficiencies have risen from around 12.7% in 2002 to 18.4% in 2018 (Barbose et al. 2019)), and this variation feeds directly in to the uncertainty distribution of the final forecast. For further discussion of the choices of variable available when applying Wright's law in the context of energy technologies, see Trancik et al. (2015).

One feature of LCOE 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 since no single cost metric can account for all costs associated with using a technology as part of a wider system — there will always be many external factors affecting costs, and these will depend on how the whole system has developed historically (i.e. the extent to which existing infrastructure has created technology lock-in). System costs are often quantified in terms of an extra component added to the LCOE of each variable renewable energy source. However, calculating this component requires very detailed modeling (e.g. OECD & NEA (2019)), which necessarily involves many more technical assumptions and layers of complexity. Instead, as described in Section 3, we have ensured that all system costs are accounted for by including (somewhat overly) generous quantities of storage capacity, dispatchable generation and grid expansion, and adding their costs separately.

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 *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 current LCOE values and model parameters for all clean 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 thus is a reasonable point of comparison to fossil fuel prices. In addition we compensate by being conservative in our choice of the mean reversion levels for fossil fuel prices.

6.1 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 BP (2019). 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 2018 was 71.31 \$/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 $c_0^1 = 11.7$ \$/GJ.

The AR(1) model is fit to data for 1975-2018 (for reasons given in Section 5.1) and the estimated parameters are $\hat{\phi} = 0.846$, $\hat{\sigma}_\epsilon = 0.252$ and $\hat{\kappa} = 0.342$. The full data set and the forecast using these estimated parameters are shown in Figure S13.

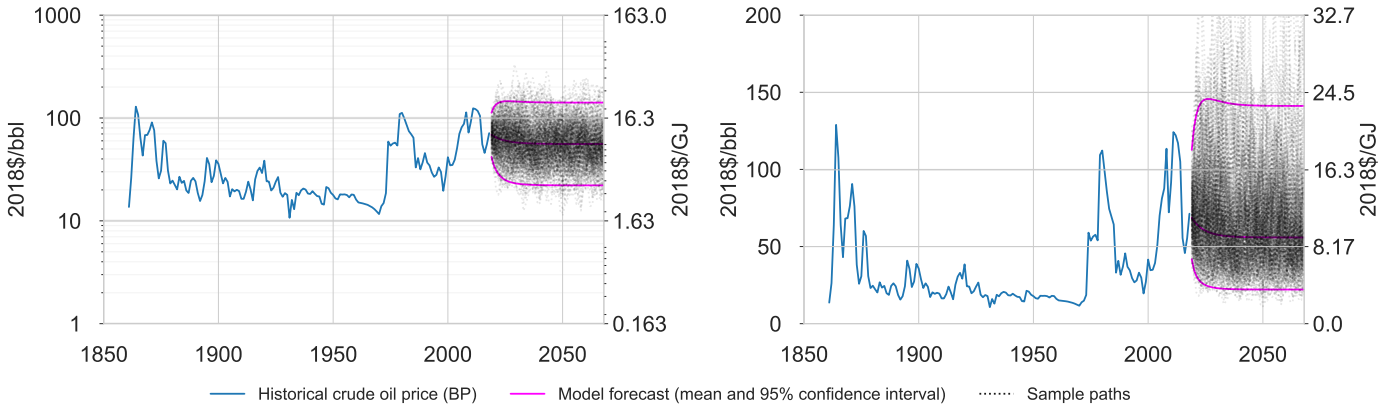

Figure S13: Crude oil price 1861-2018, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

The oil cost shown here is the cost of crude oil, but our energy model considers bulk flows of refined oil. Oil refining costs around 7.5 \$/bbl, or about 1.0 \$/GJ (see Table S20), and is roughly constant over time, so when calculating the cost of refined oil in the model, we add this constant value to the cost of crude oil given by the AR(1) process.

6.2 Coal

Annual coal price data since 1882, collated from various sources, is provided by McNerney et al. (2011). 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 coal (FRED 2019). These are shown in Figure S14.

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 applications (Vaillancourt 2014). 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 Section 6.4, 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: $\hat{\phi} = 0.954$, $\hat{\sigma}_\epsilon = 0.090$ and $\hat{\kappa} = 0.035$. We initialize the model with a current cost $c_0^2 = 2.18$ \$/GJ (see Table S20).

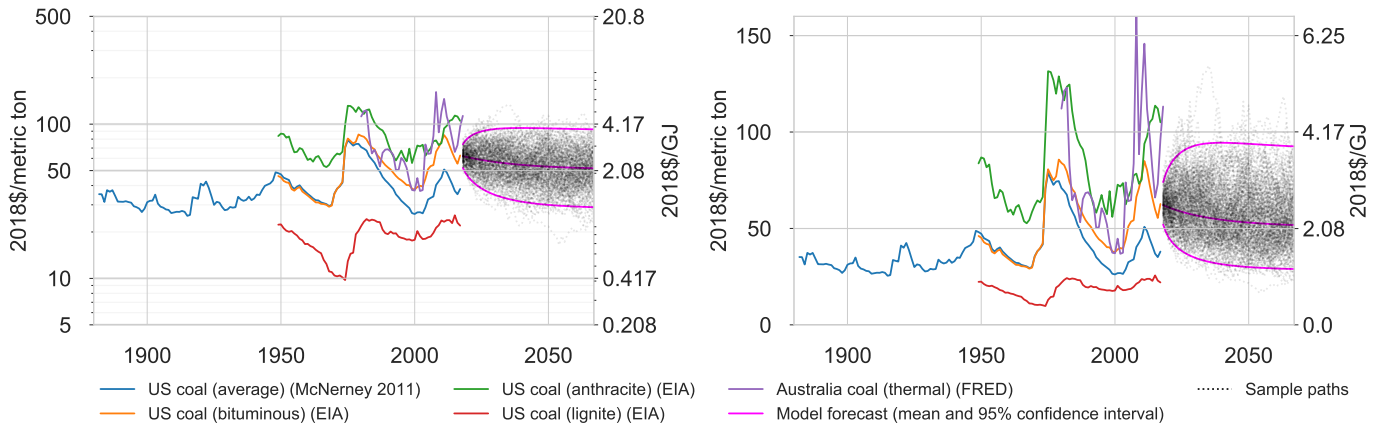

Figure S14: Coal price data 1882-2018, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

6.3 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 FRED. Figure S15 shows the data (1 MMBtu is approximately equal to 1.05506 GJ). As with coal, the price of gas varies depending on location (which we ignore). The real

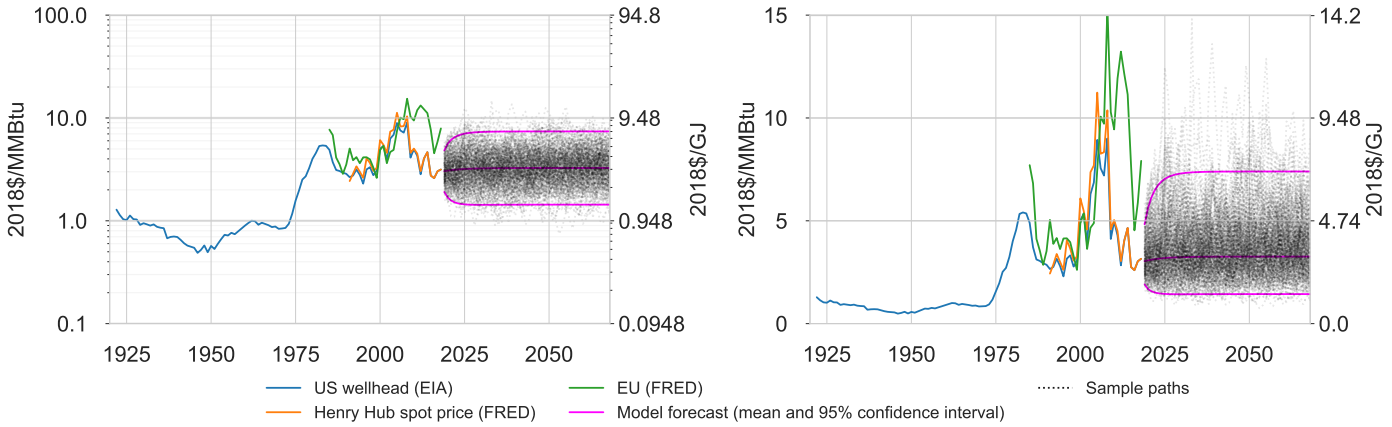

Figure S15: Gas price data 1922-2018, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

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 1975-2018 (for reasons given in Section 5.1). The estimated parameters are: $\hat{\phi} = 0.824$, $\hat{\sigma}_\epsilon = 0.237$ and $\hat{\kappa} = 0.207$. We take the initial cost to be $c_0^3 = 3$ \$/GJ (see Table S20).

6.4 Coal electricity

Coal electricity generation data since 1960 is available from the IEA, McNerney et al. (2011) 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. (2011), with Lazard and the EIA providing more recent data also. The time series and experience curve are shown in Figure S16. Underlying coal prices from Section 6.2 are also shown on the time series plot, to highlight the connection between the two series since 1960.

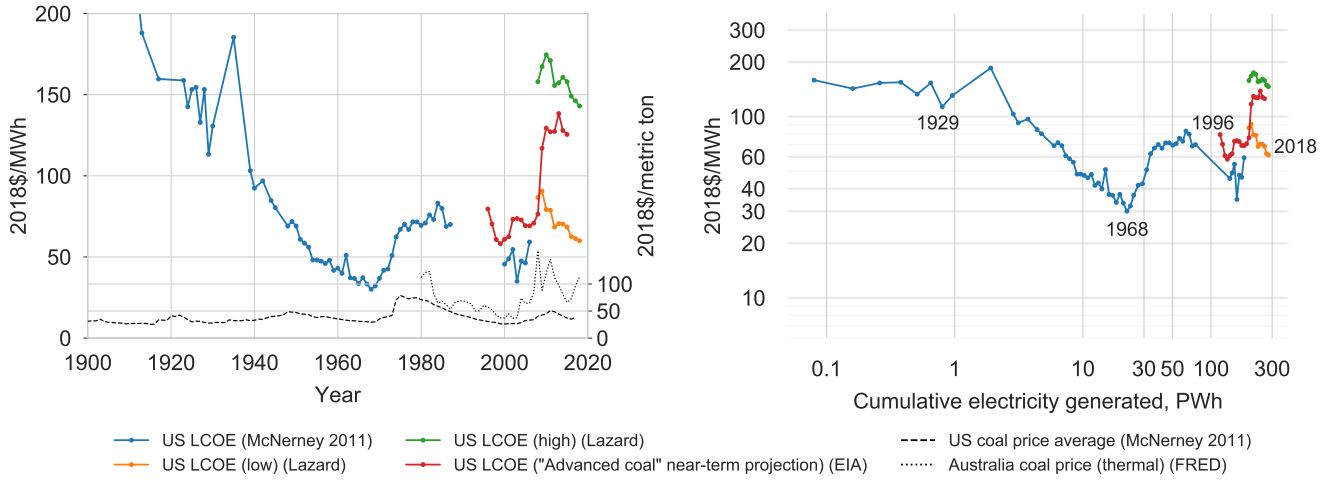

Figure S16: US coal electricity LCOE time series (L) and experience curves (R). Two coal price time series from Figure S14 are also shown on (L). Generation data since 1960 is from the IEA, McNerney et al. (2011), BP, and before this it is our own estimate.

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 2018$/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 S17 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.[16]

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

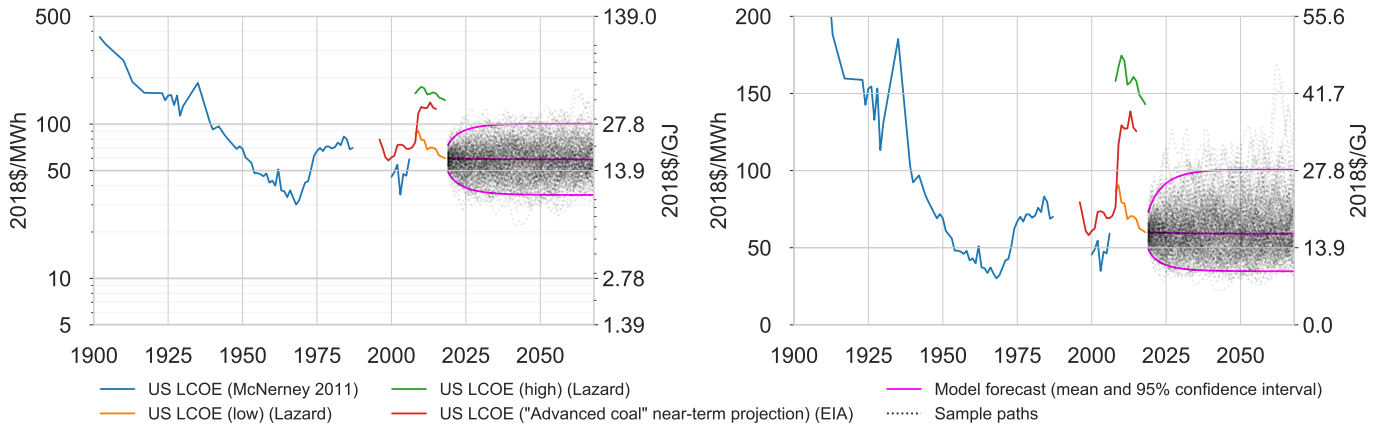

Figure S17: Coal LCOE data 1902-2018, AR(1) forecast plus 95% confidence interval and 200 sample paths; on a log scale (L) and a regular scale (R).

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

6.5 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 (Fouquet 2008, Smil 2016). 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 & Cornland (2002). 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 utilization). 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 costs (Seebregts 2010), 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 S18. The underlying gas price is also shown.

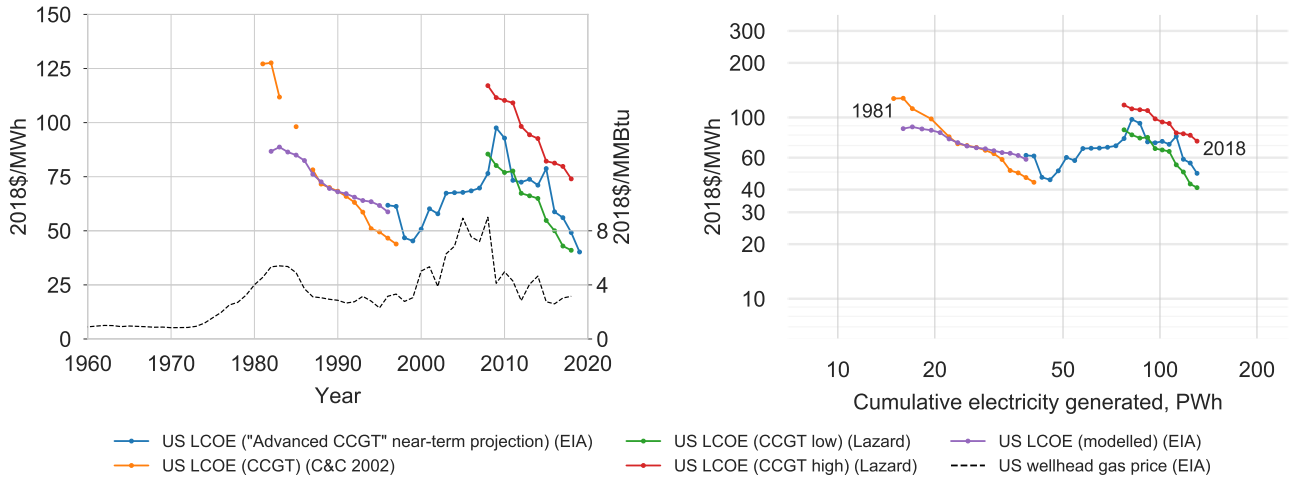

Figure S18: Gas electricity LCOE time series (L) and experience curves (R). A gas price time series from Figure S15 is also shown on (L). Generation data since 1973 is from the IEA and BP.

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 S19, and the parameters estimated are: $\hat{\phi} = 0.827042$, $\hat{\sigma}_\eta = 0.130628$ and $\hat{\kappa} = 0.48506$.

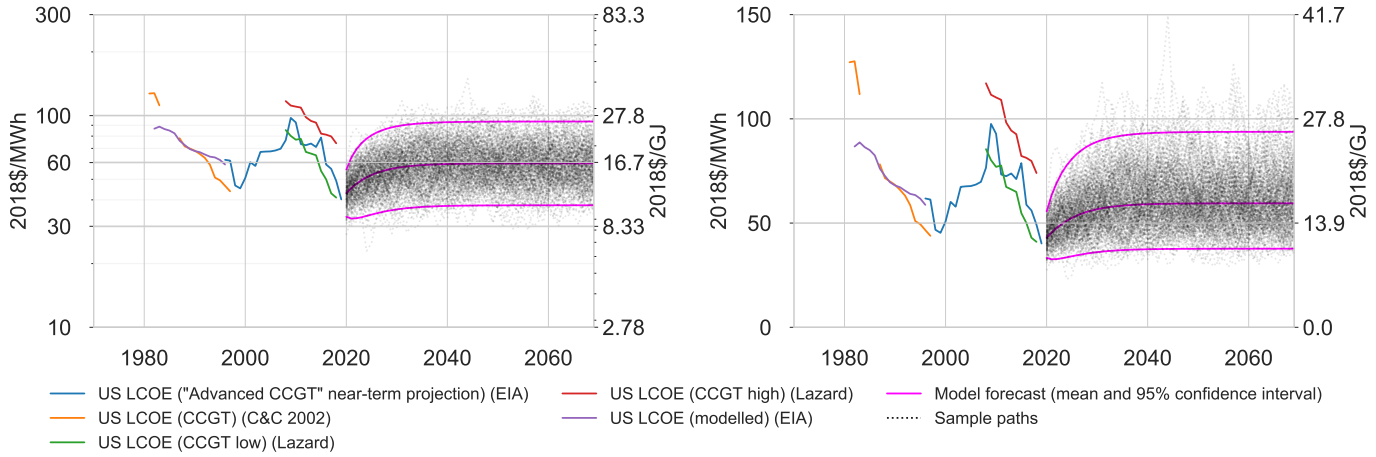

Figure S19: Gas LCOE data 1996-2019. Data: US EIA. Also shown is the AR(1) forecast plus 95% confidence interval, and 200 sample paths.

6.6 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, specific investment cost etc.). This is the investment per kW of capacity that would be required to install a nuclear reactor very quickly (e.g. "overnight") — the cost of instantly buying and installing the required technology. Overnight costs generally differ a lot from total costs 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. (2016), 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 the reasons why they can give a very misleading impression of the total costs of nuclear energy were detailed by Koomey et al. (2017) and Gilbert et al. (2017).

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 overruns (Sovacool et al. 2014), 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. Grubler (2010) analysed capacity investment costs of the French nuclear program and Matsuo & Nei (2019) 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 & Hultman (2007), while those in France were analysed by Boccard (2014) (based on data from the Court of Audit (2012)), 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 S20. In addition, Figure S21 shows all the capacity cost data we could find (OCCs and investment costs).

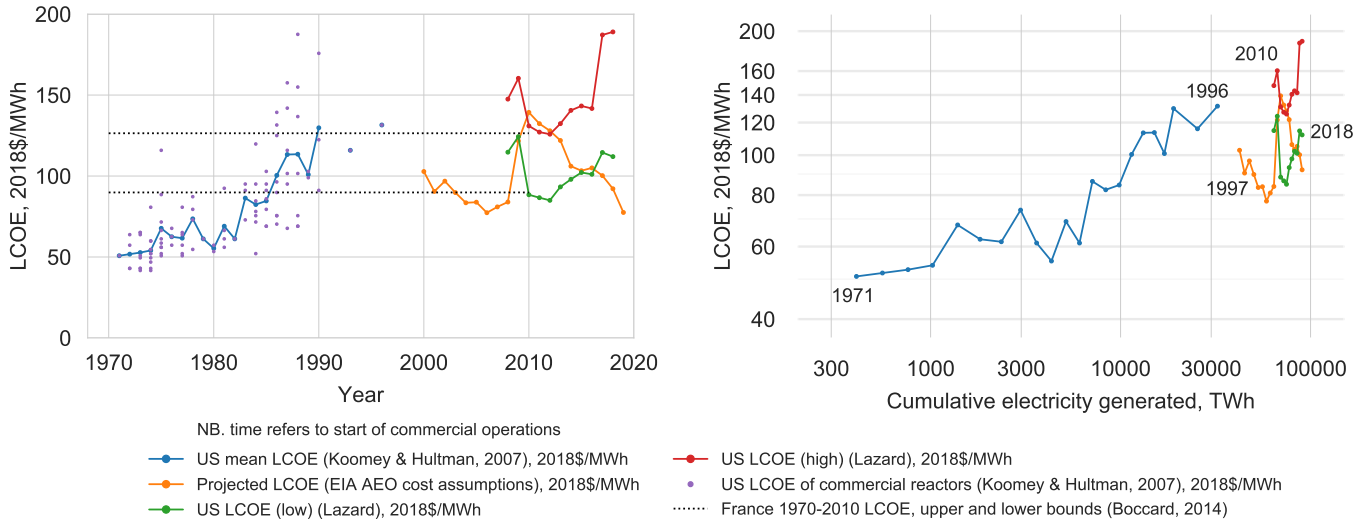

Figure S20: Nuclear power LCOE time series data (L) and experience curves (R).

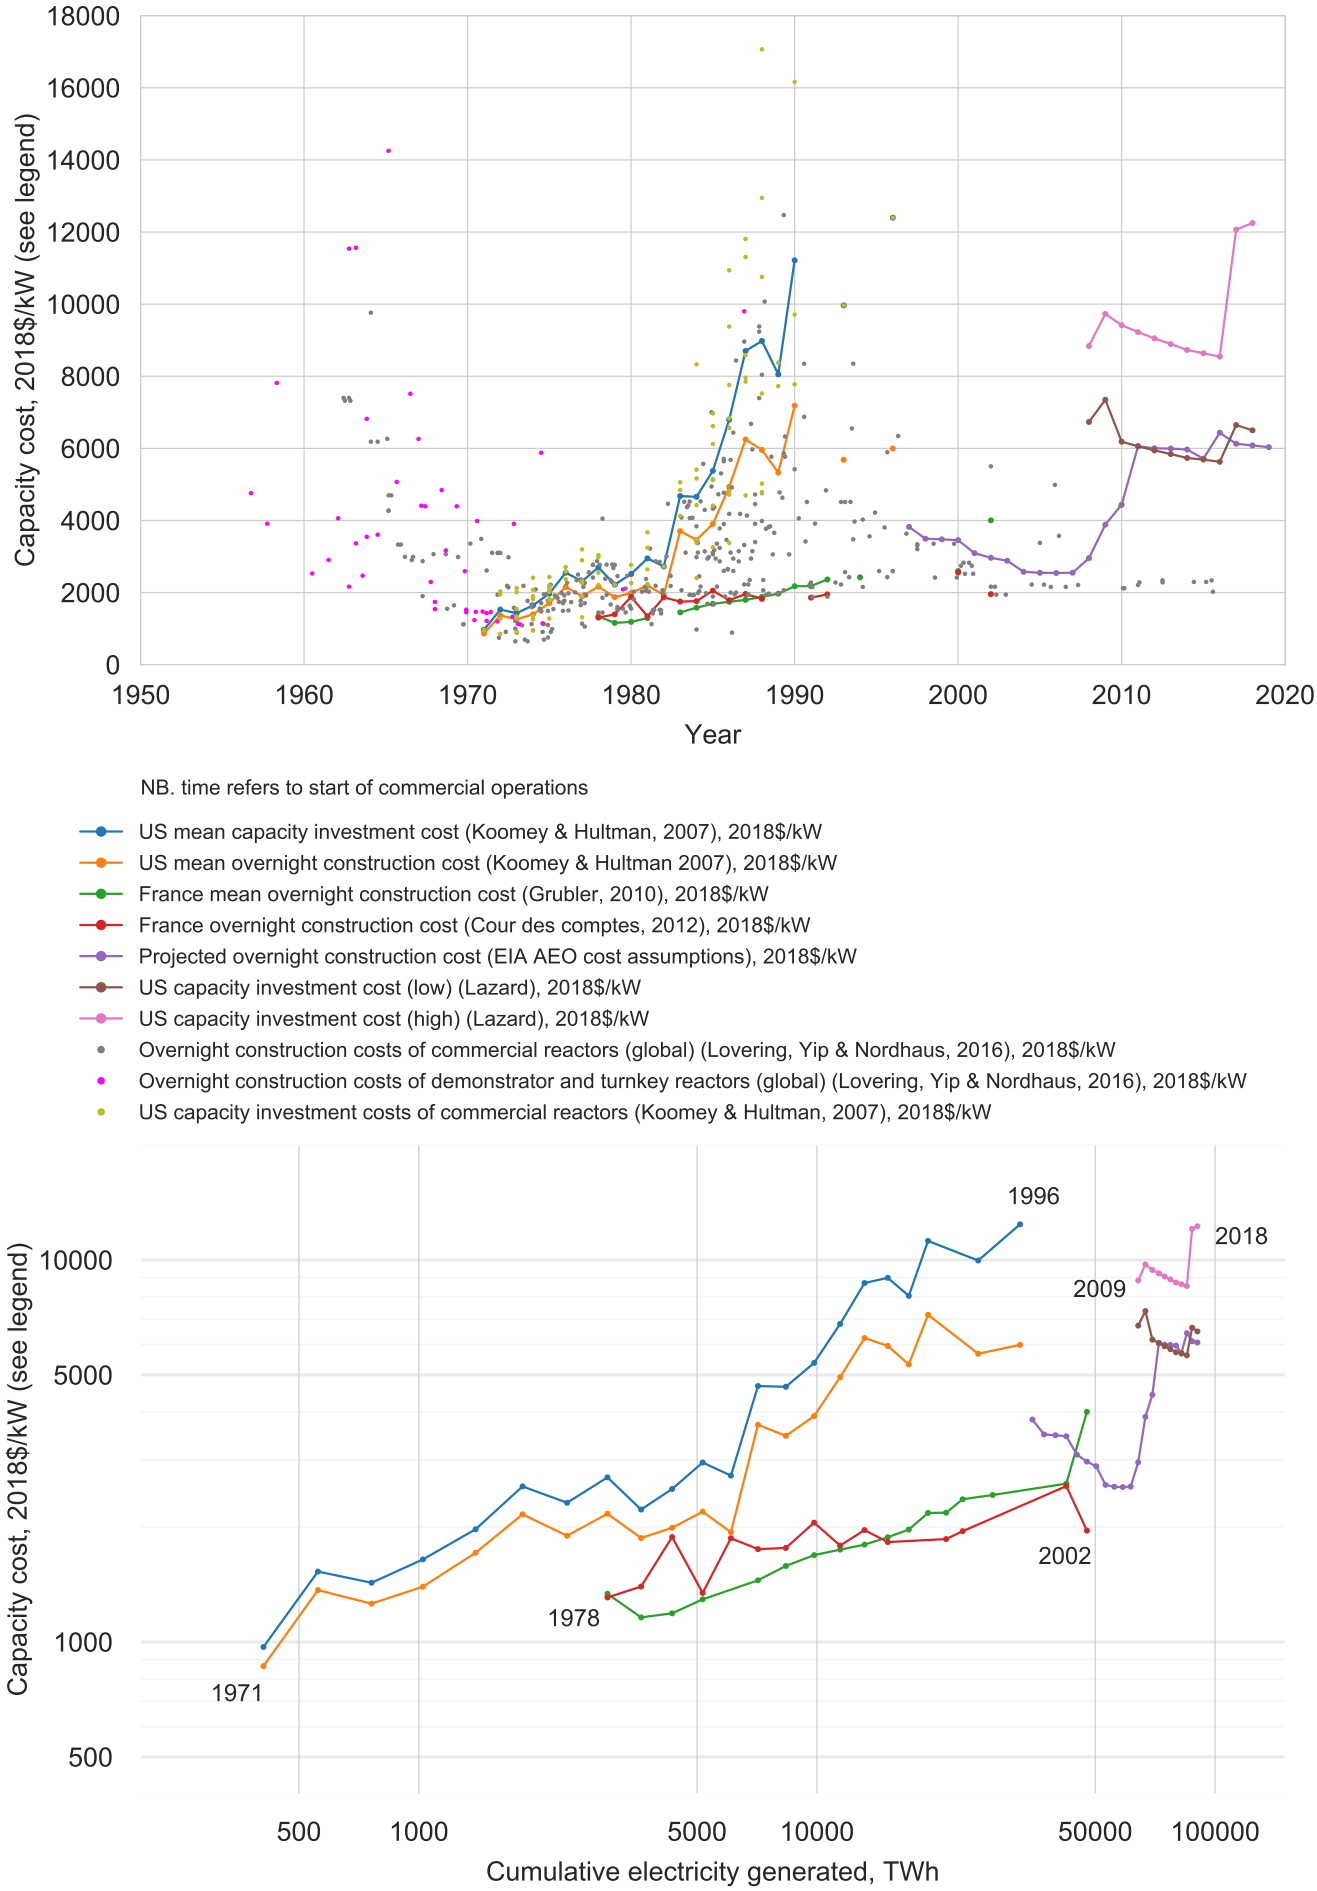

Figure S21: Nuclear power capacity investment cost data.

Long-term trends in nuclear power are much discussed, but empirically it is clear that on average, globally, costs have increased, even though a great deal of experience in generating electricity in large, stationary power plants has been accumulated. Note that in addition, substantial experience has been gained in designing, building and operating small, portable nuclear reactors, as these have been used widely in ships and submarines.

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 S20) 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 S22 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 summarized in Table S10.

	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 S10: Nuclear parameters.

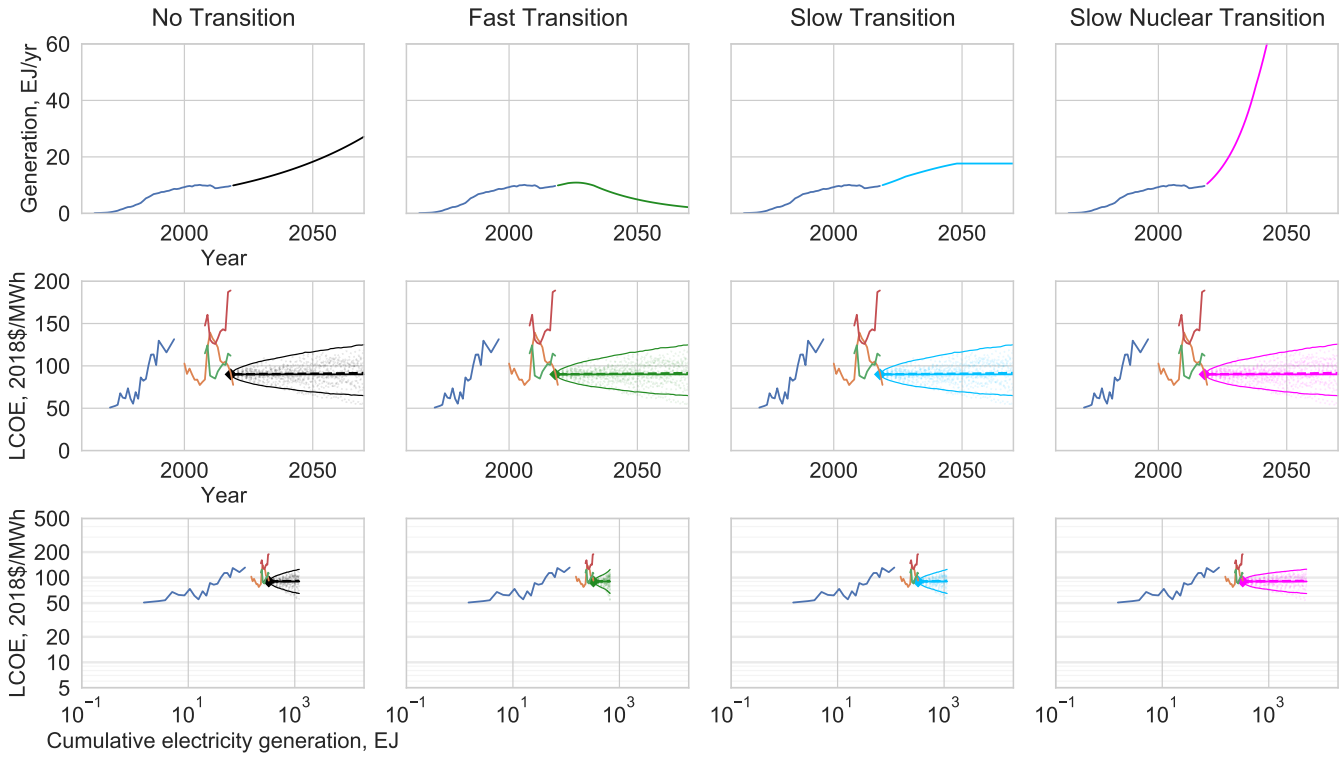

Figure S22: Nuclear electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S20.

6.7 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 IRENA (2019) 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 S23 and S24.

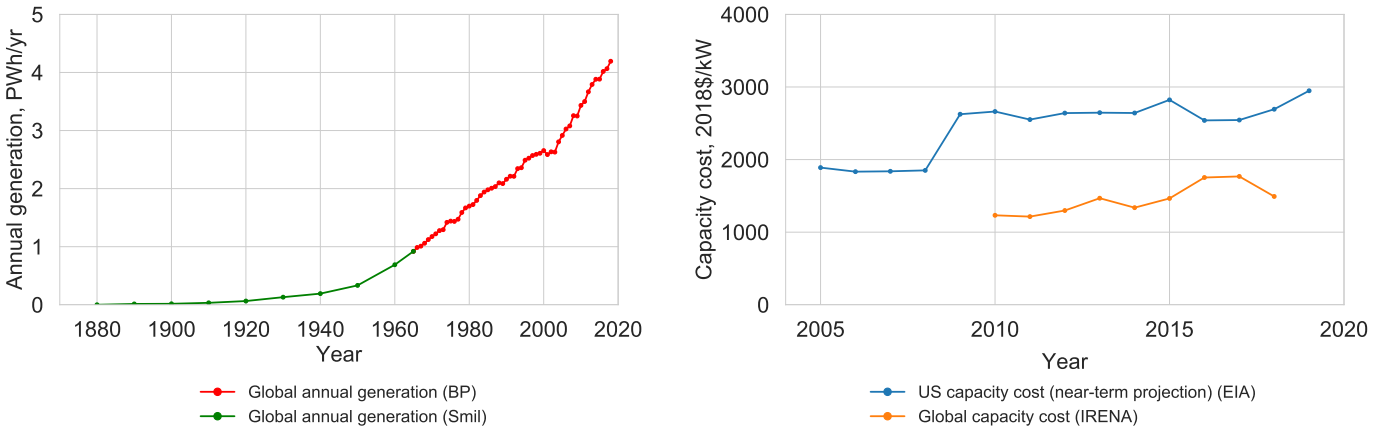

Figure S23: Hydropower annual generation (L) and capacity cost time series (R).

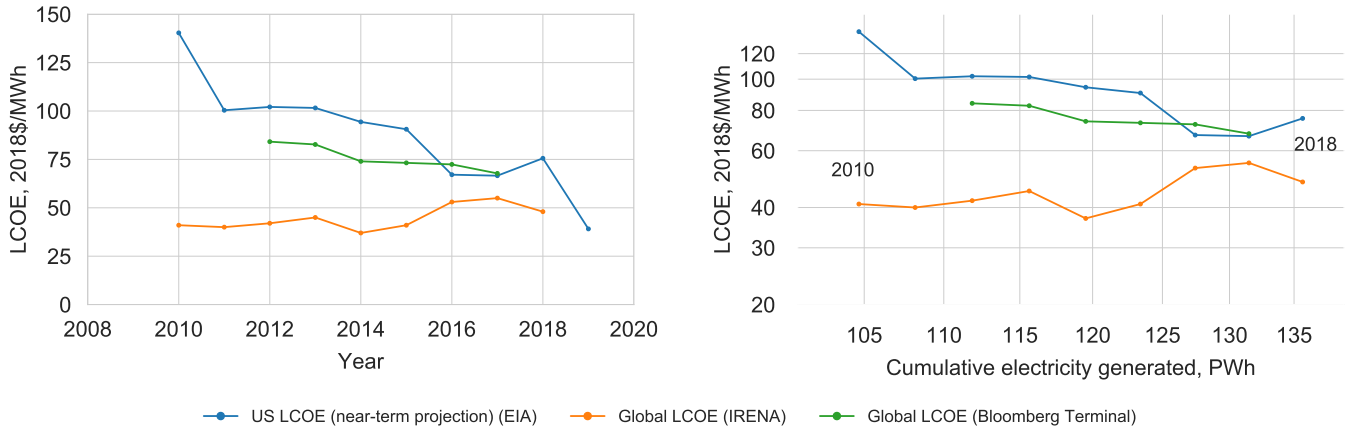

Figure S24: Hydropower LCOE time series (L) and experience curves (R).

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 modeling 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 common (Ansar et al. 2014).

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 S25 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 summarized in Table S11.

	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 S11: Hydropower parameters.

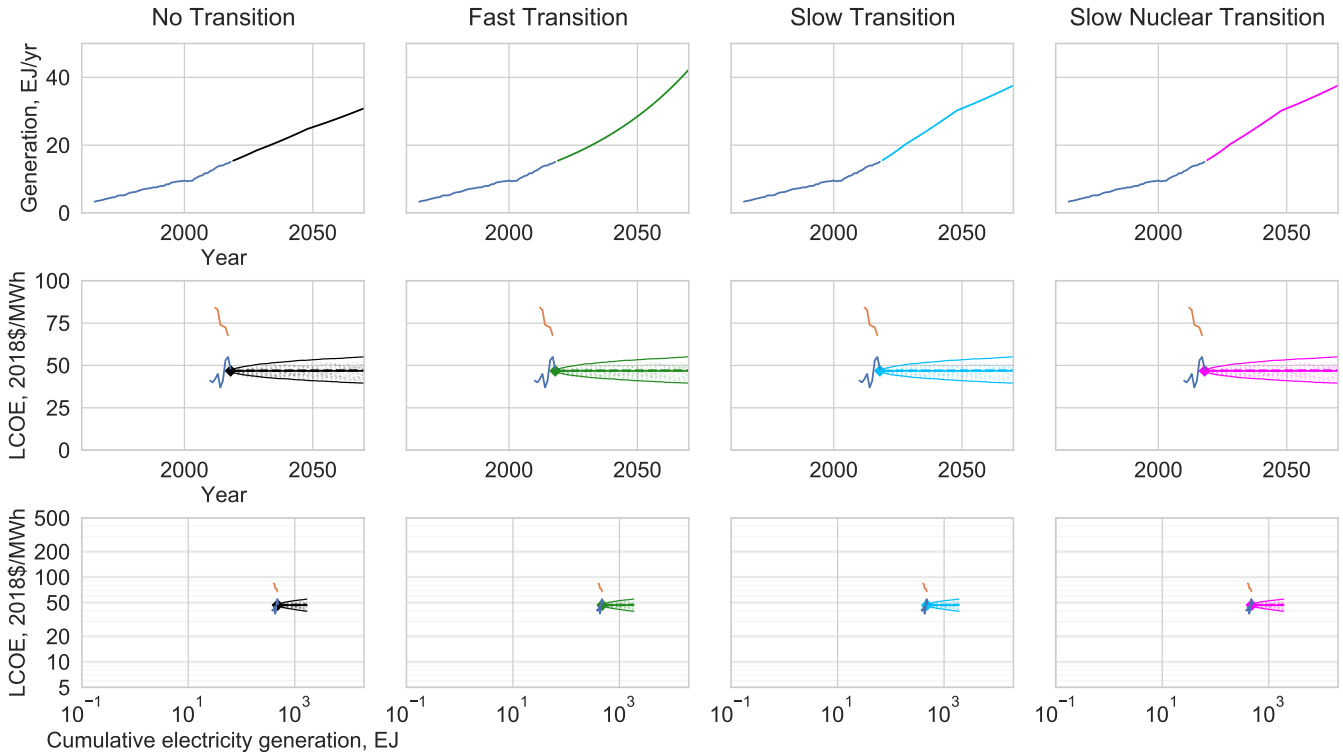

Figure S25: Hydropower electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S24.

6.8 Biopower electricity

For biopower generation data we use the Biofuels and Waste category in the IEA's Energy Technology Perspectives 2017 report. 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. (2006), 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 S26.

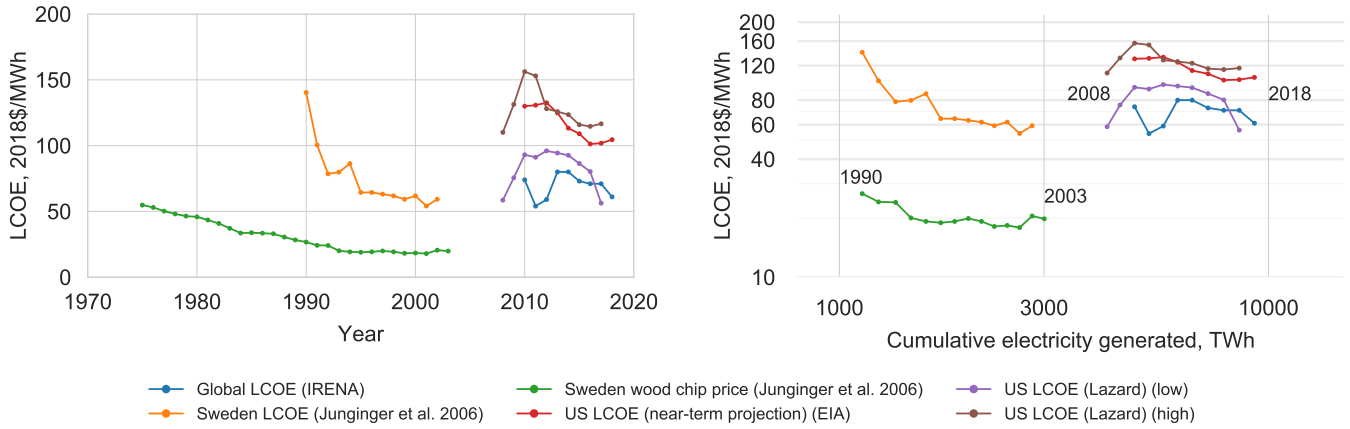

Figure S26: Biopower LCOE time series (L) and experience curves (R).

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% (Rubin et al. 2015 - A review of learning rates for electricity supply technologies), 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 S27 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 S12.

	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 S12: Biopower parameters.

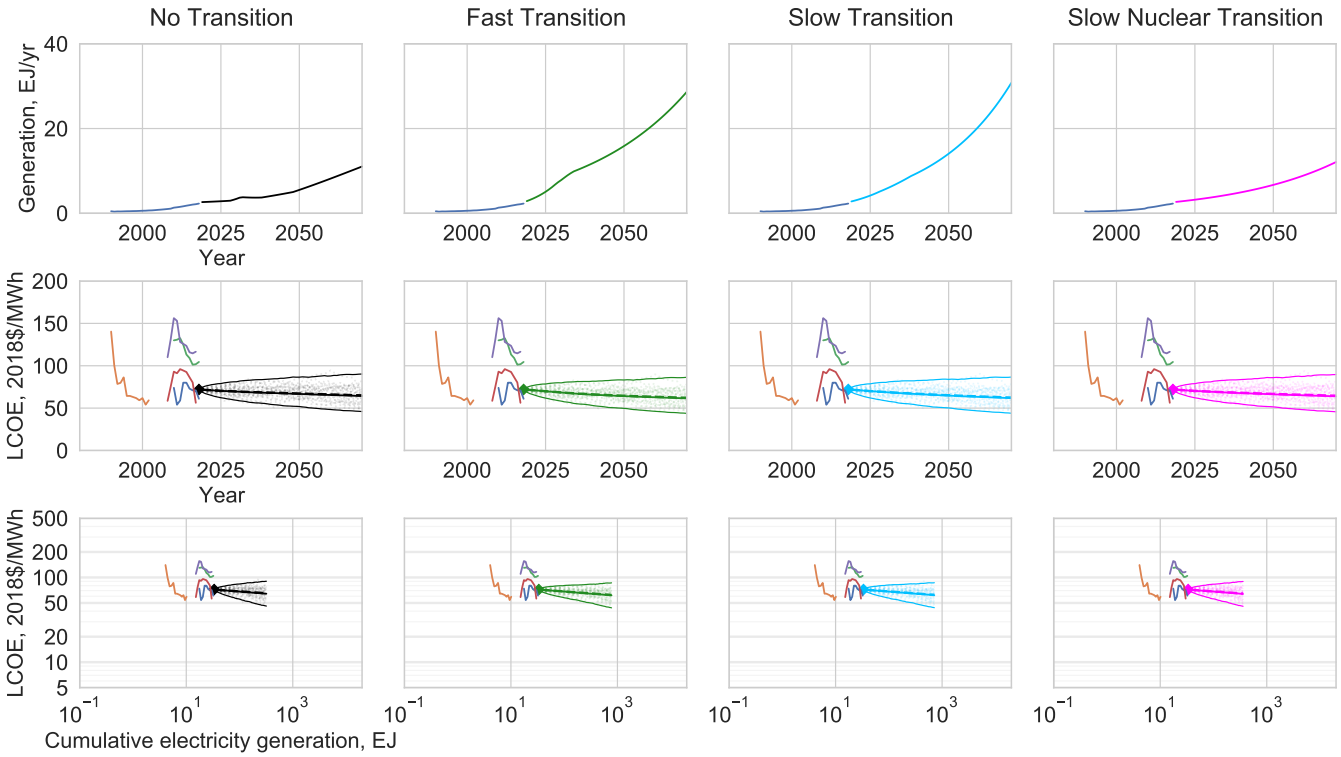

Figure S27: Biopower electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S26.

6.9 Wind electricity

We take global annual wind electricity generation data from BP (2019). We compared this with several other sources such as IRENA and the IEA to verify broad agreement. We have 11 different cost series for the LCOE of onshore wind, from Wiser et al. (2016), Schilling & Esmundo (2009), Lazard (2019*a*), IRENA (2018*b*), IRENA (2019) and Bloomberg Terminal. Time series, experience curves and data sources are shown in Figure S28. The parameters obtained by fitting the FDWL model, plus the length of series, are shown in Table S13. Two of the 11 cost series are excluded due to the data being mostly non-sequential. There are two

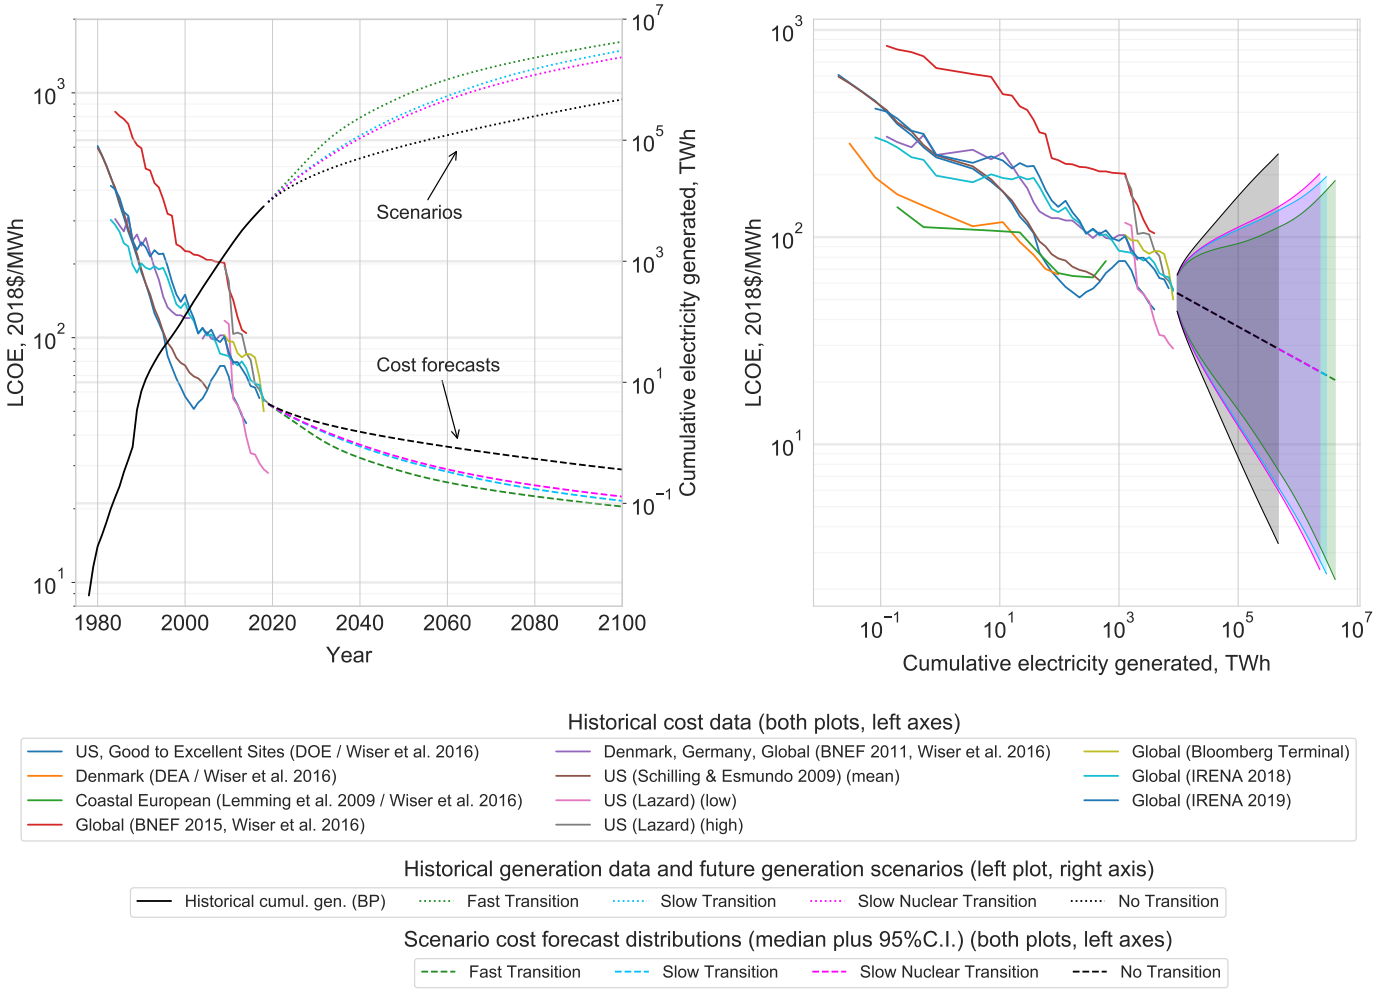

Figure S28: Wind LCOE time series and forecast (L, left axis). Historical generation and future scenarios (L, right axis). Historical experience curves and cost forecasts in each scenario (R).

very short series, giving experience exponents very different to the others, as the cost reductions since 2010 have been especially sharp; these are not representative of the long-term trend and are excluded from further analysis. However, we have chained together the BNEF 2015 series with the final four points of the Bloomberg Terminal series to bring the former right up to 2018, thus hopefully improving the representativeness of both of the individual series. We take $\hat{\omega} = 0.158$, $\hat{\sigma}_\omega = 0.045$ and $\hat{\sigma}_\eta = 0.103$. Figure S29 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 $c_0^9 = 55$ \$/MWh. Parameter choices are summarised in Table S14.

Data source	$\hat{\omega}$	$\hat{\sigma}_\omega$	$\hat{\sigma}_\eta$	Start year	Final year	Series length	Current LCOE
Original data							
US DOE (Wiser et al. 2016)	0.183	0.032	0.077	1980	2014	35	
BNEF 2015 (Wiser et al. 2016)	0.137	0.037	0.084	1984	2014	31	
BNEF 2011 (Wiser et al. 2016)	0.069	0.044	0.097	1984	2011	26	
Schilling & Esmundo (2009) (mean)	0.184	0.020	0.048	1980	2005	26	
Lazard LCOE Analysis Low	0.793	0.332	0.207	2009	2018	10	28 $/MWh
Lazard LCOE Analysis High	0.691	0.241	0.150	2009	2018	10	54 $/MWh
Bloomberg Terminal / BNEF	0.349	0.187	0.117	2009	2018	10	50 $/MWh
IRENA (2018*b*)	0.126	0.034	0.079	1983	2017	35	56.6 $/MWh
IRENA (2019)	0.101	0.029	0.068	1983	2018	36	55 $/MWh
Chained data							
BNEF 2015 + Bloomberg Term.	0.158	0.045	0.103	1984	2018	35	50 $/MWh

Table S13: Wind data sources and the associated estimated parameters. (The series length may be less than the difference between the start and final years if data is missing for some years.)

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.158$	
Experience exponent standard error:	$\hat{\sigma}_\omega = 0.045$	
Noise standard deviation:	$\hat{\sigma}_\eta = 0.103$	
Current cost:	$c_0^9 = 55$ $/MWh	$c_0^9 = 50$ $/MWh
Lifespan:	$L^9 = 30$	

Table S14: Wind parameters.

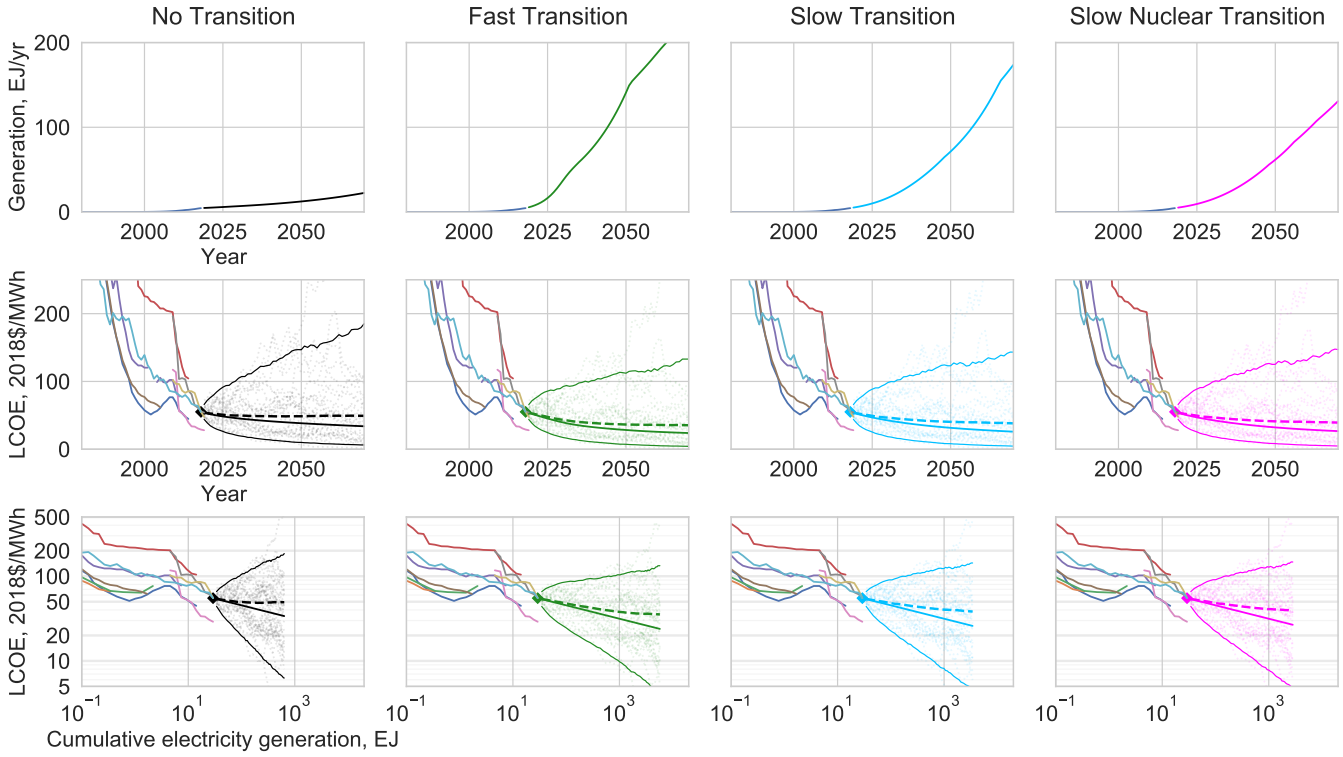

Figure S29: Wind electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S28.

6.10 Solar PV electricity

Solar PV generation data for 1990-2014 is from IEA (2017) and after that from BP (2019); we did not find any data before this. LCOE data is from a wide variety of sources. Historical time series and experience curves are shown in Figure S30, along with generation in each of the four scenarios, plus the implied cost forecast distribution in each case. FDWL is a good model for this technology. Experience curve parameters are shown

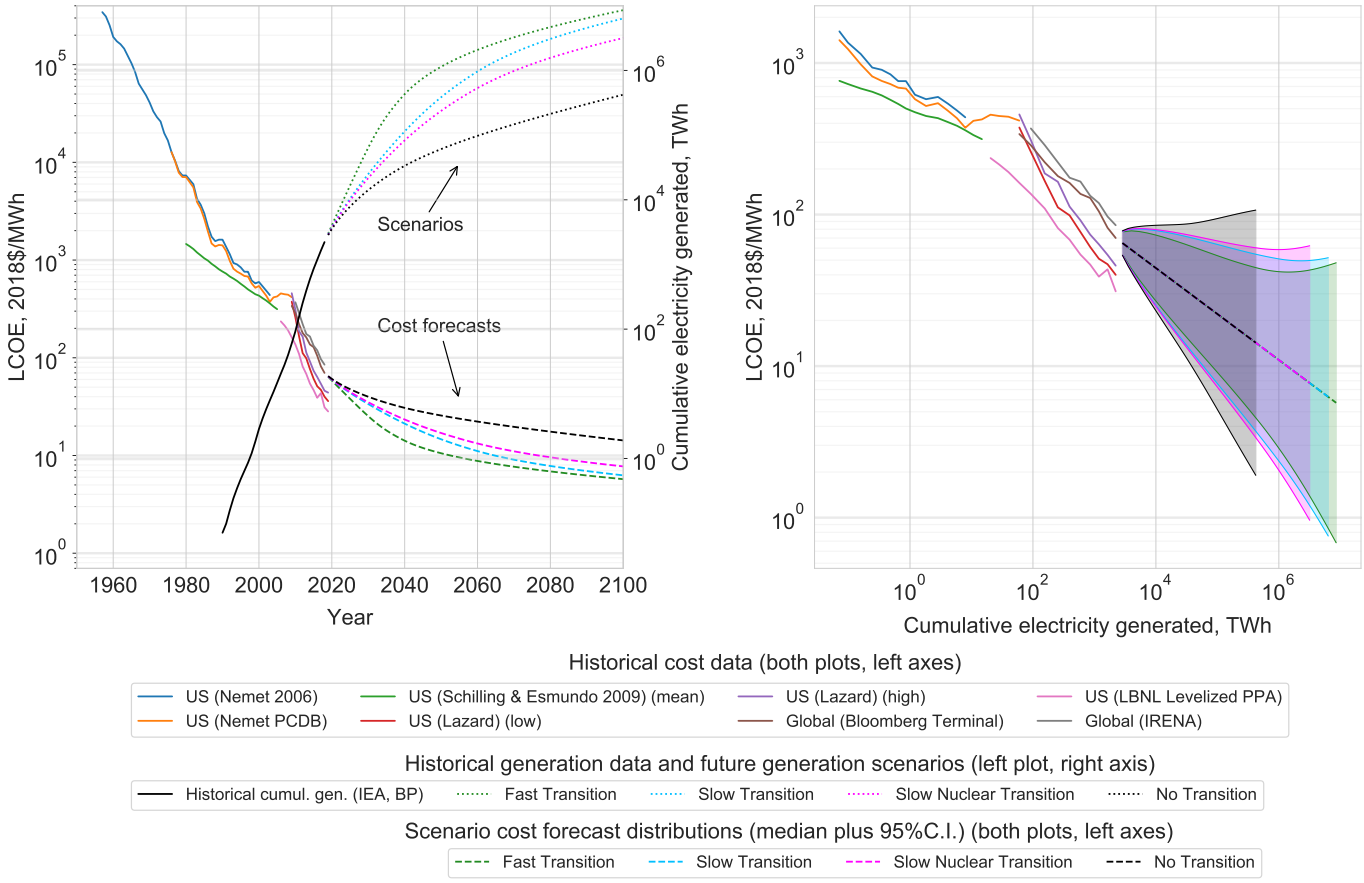

Figure S30: Solar PV LCOE time series and forecast (L, left axis). Historical generation and future scenarios (L, right axis). Historical experience curves and cost forecasts in each scenario (R).

in Table S15.

We take $\hat{\omega} = 0.303$, $\hat{\sigma}_\omega = 0.047$ and $\hat{\sigma}_\eta = 0.093$. Figure S31 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 an initial cost value of $c_0^{10} = 70$ \$/MWh. Parameter choices are summarised in Table S16.

Data source	$\hat{\omega}$	$\hat{\sigma}_\omega$	$\hat{\sigma}_\eta$	Start year	Final year	Series length	Current LCOE
Original data							
Nemet (2006)	0.273	0.057	0.075	1990	2003	14	
Nemet PCDB	0.192	0.055	0.086	1990	2009	20	
Schilling & Esmundo (2009) (mean)	0.160	0.012	0.017	1990	2005	16	
Lazard LCOE Analysis (Low)	0.643	0.077	0.094	2009	2018	10	40 $/MWh
Lazard LCOE Analysis (High)	0.652	0.090	0.110	2009	2018	10	46 $/MWh
Bloomberg Terminal / BNEF	0.430	0.054	0.065	2009	2018	10	70 $/MWh
LBNL Levelized PPA	0.436	0.076	0.104	2006	2018	13	28.2 $/MWh
IRENA (2019)	0.465	0.053	0.061	2010	2018	9	85 $/MWh
Chained data							
Nemet PCDB and IRENA (2019)	0.287	0.046	0.091	1990	2018	29	85 $/MWh
Nemet PCDB and Bloomberg Terminal	0.303	0.047	0.093	1990	2018	29	70 $/MWh

Table S15: Solar PV data sources and the associated estimated parameters.

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.303$	
Experience exponent standard error:	$\hat{\sigma}_\omega = 0.047$	
Noise standard deviation:	$\hat{\sigma}_\eta = 0.093$	
Current cost:	$c_0^{10} = 70$ $/MWh	$c_0^{10} = 50$ $/MWh
Lifespan:	$L^{10} = 30$	

Table S16: Solar parameters.

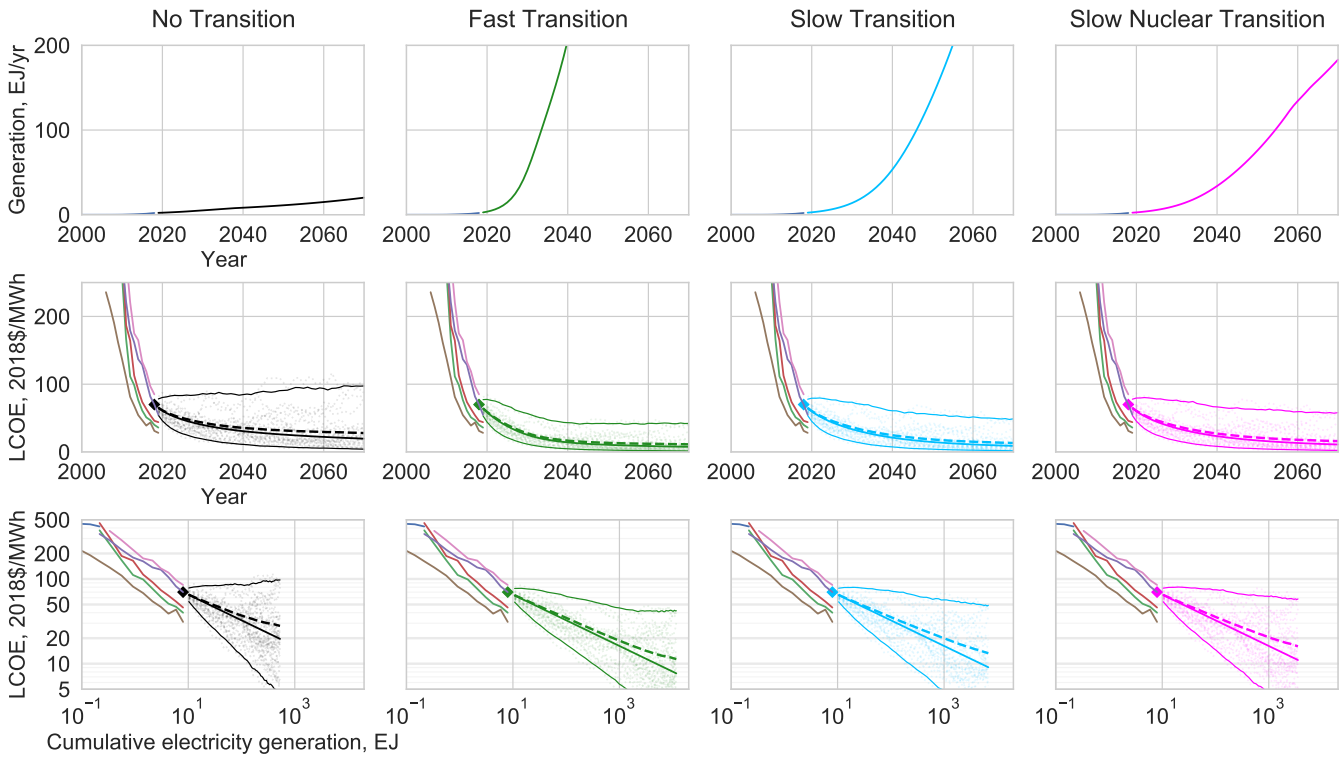

Figure S31: Solar PV electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S30.

6.11 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 IEA (2018), Schmidt et al. (2017), plus reports and presentations freely available on the internet from Citibank (2016), Pillot (2018), BNEF (2019), Clark (2019) and Lazard (2019*b*). For Va-redox flow batteries, data comes from Schmidt et al. (2017) and Lazard (2019*b*). Figure S32 shows the time series and experience curves.

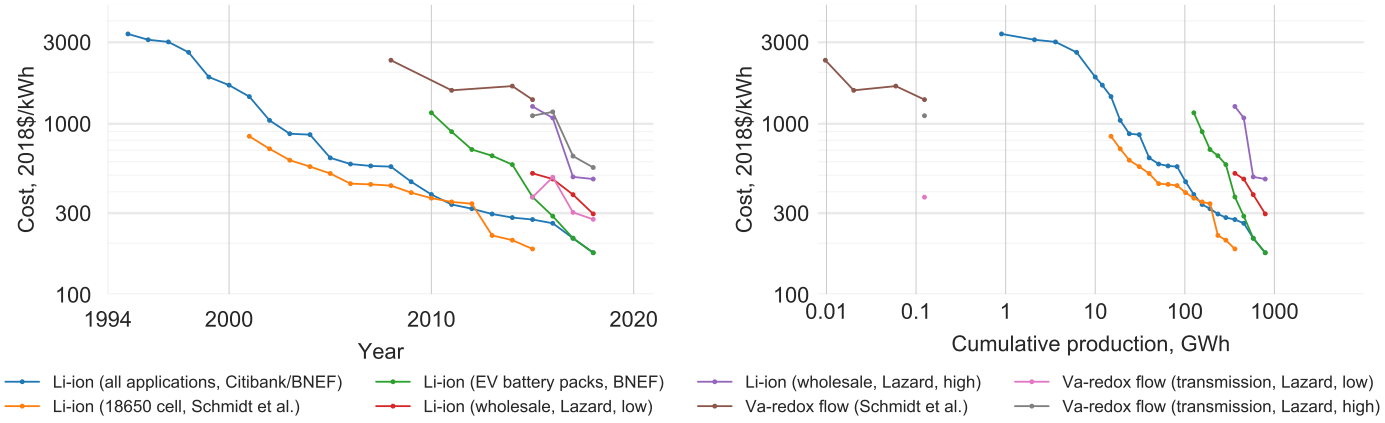

Figure S32: Battery cost time series (L) and experience curves (R).

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 technology lifetime to be $L^{11} = 10$ years. Fu et al. (2018) report utility-scale Li-ion storage costs of 380 \$/kWh, but pack costs have fallen several tens of dollars per kWh further since then, so we take $c_0^{11} = 360$ \$/kWh as the initial cost value. Figure S33 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, Schmidt, Staffell & Kammen (2020), while the role of battery technologies in EVs in particular is considered in Nykvist & Nilsson (2015), Kittner et al. (2017) and Kittner, Tsiropoulos, Tarvydas, Schmidt, Staffell & Kammen (2020). The progress and innovation rates found in these studies are consistent with our modeling.

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. (2017) reported a learning rate of 11±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 very close to the value for wind electricity (0.158) so we assume that the other two parameters are similar also, and use these: we take $\hat{\omega} = 0.168$, $\hat{\sigma}_\omega = 0.045$ and $\hat{\sigma}_\eta = 0.103$. 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 2019 (Lazard 2019*b*) 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. Figure S34 shows production in each scenario with the associated cost forecasts given these parameters. Parameter choices are summarized in Table S17.

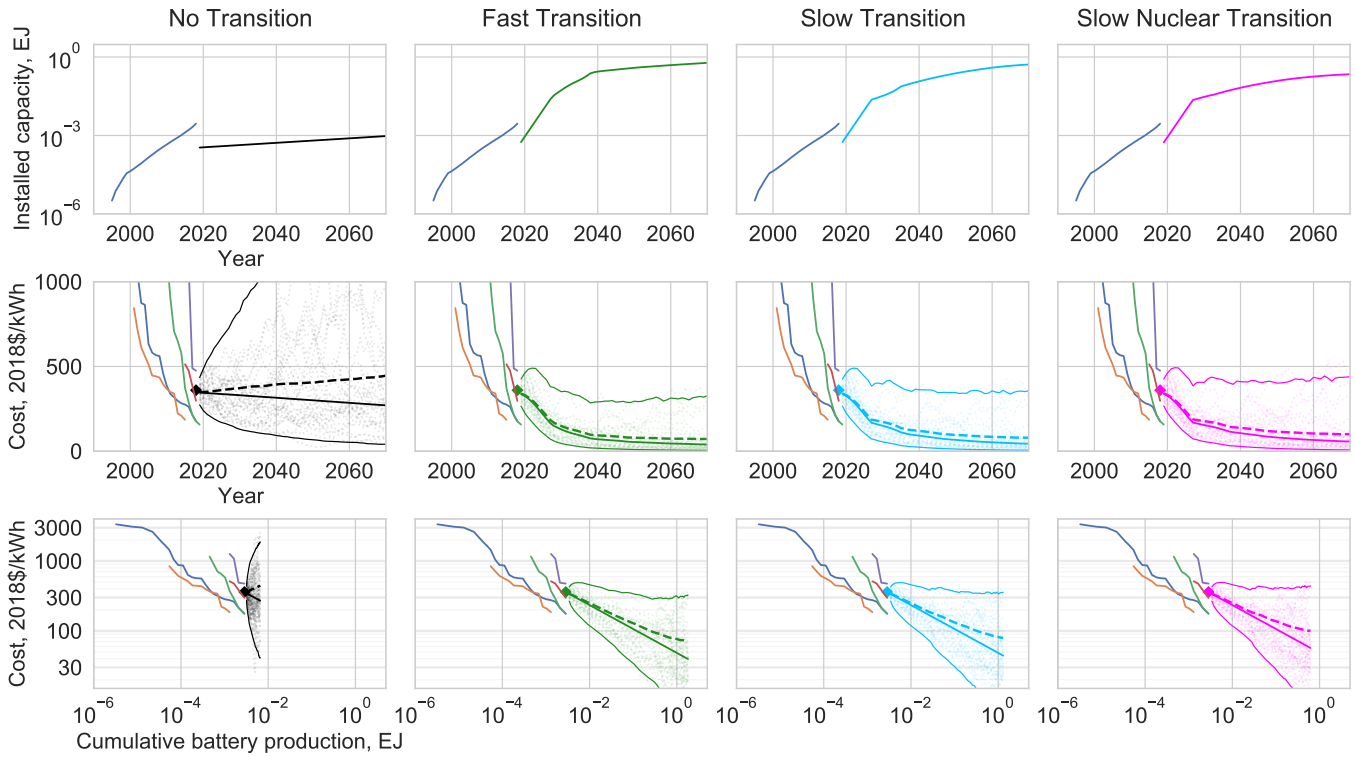

Figure S33: Daily battery storage capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S32. 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.

	Main case (Li-ion)	Main case (flow battery)
Experience exponent:	$\hat{\omega} = 0.342$	$\hat{\omega} = 0.168$
Experience exponent standard error:	$\hat{\sigma}_\omega = 0.073$	$\hat{\sigma}_\omega = 0.045$
Noise standard deviation:	$\hat{\sigma}_\eta = 0.116$	$\hat{\sigma}_\eta = 0.103$
Current cost:	$c_0^{11} = 360$ \$/kWh	$c_0^{12} = 400$ \$/kWh
Lifespan:	$L^{11} = 10$	$L^{12} = 20$

Table S17: Battery parameters.

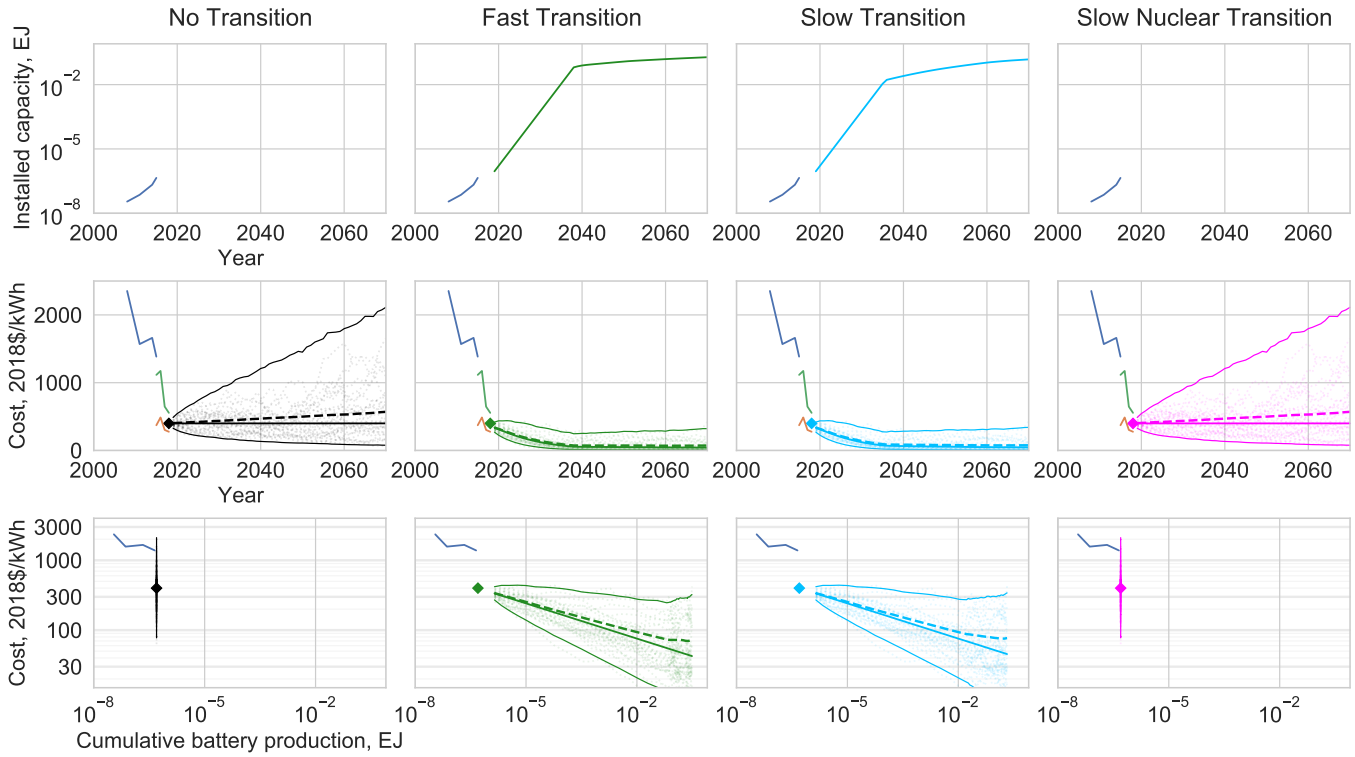

Figure S34: Multi-day battery storage capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S32.

6.12 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, for calibrating the FDWL model and forecasting costs.

Current global P2X fuels production is around 115 Mt (million tonnes) (IEA 2019*b*). 1 Mt H_2 is 1×10^9 kg H_2, 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 H_2 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. The final production method is electrolysis of water, making up just a 2-4% of the total. Despite their current dominance, decarbonization of fossil fuel based methods 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 technologies are very closely linked. Indeed some current electrolyzers can be run in reverse, providing fuel cell functionality. We assume that electrolyzers used for electric grid backup in our model will have this capability, running in one direction when excess renewable electricity is available and in the other when there is a scarcity. Others electrolyzers may be dedicated to hydrogen production for other purposes (e.g. industrial uses) and not require this functionality. There are three established types of electrolysis and electrolyzer, with different characteristics and potential for future development: alkaline electrolysis (AEL), polymer electrolyte membrane electrolysis (PEM) and solid oxide electrolysis cells (SOEC) (see e.g. Buttler & Spliethoff (2018), Carmo et al. (2013), IRENA (2018*a*) 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 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, and 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 (Schoots et al. 2008, Schmidt et al. 2017). Early installations were placed near hydropower facilities to take advantage of cheap, reliable electricity, but this practice stopped as cheaper fossil-based methods took over. Both PEM and SOEC have much lower cumulative installed capacities, probably around 140 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 seems likely 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, with reversible PEM electrolyzers also 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 (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 suggest that they perform different functions and are part of different technology clusters — they are often not substitutes in applications. 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 very 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 referenced: Schoots et al. (2008), Schmidt et al. (2017) and the accompanying dataset, IEA (2019*b*) and the accompanying dataset, Buttler & Spliethoff (2018), Staffell et al. (2019). The first two of these consider the full history of AEL applications, while all the others consider electrolyzer capacity "for energy purposes" only, i.e. PtG systems — this is an important difference. Buttler & Spliethoff (2018) provide some capacity data from the 1990s, so to construct cumulative PtG capacity we use this data initially, then the IEA data subsequently. Note however that the Buttler & Spliethoff (2018) data is for flexible power-to-gas *and* power-to-liquid systems; we ignore the distinction and use all the data.

For cost data our primary source was the dataset accompanying Glenk & Reichelstein (2019) (abbreviated "G&R 2019" henceforth), though we also referred to the Schmidt et al. (2017) dataset. The former refers to PtG applications only, while the latter considers all AEL applications. In all cases costs refer to electrolyzer system cost. A full PtG system costs about 15% more (see G&R 2019), which is typically well within the range of cost estimates anyway, so the difference is not important. We plot the data in Figures S35-S37.

Figure S35 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. The IEA report 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 places.

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 so much lower.[17] Since we assume PEM systems have no other experience besides PtG systems, this perspective gives a fairer comparison between the two technologies. Figures S36 and S37 show the PtG plots for AEL and PEM systems.

	Main case	Side case
Experience exponent:	$\hat{\omega} = 0.194$	$\hat{\omega} = 0.1$
Experience exponent standard error:	$\hat{\sigma}_\omega = 0.090$	$\hat{\sigma}_\omega = 0.04$
Noise standard deviation:	$\hat{\sigma}_\eta = 0.233$	$\hat{\sigma}_\eta = 0.08$
Current cost:	$c_0^{13} = 1100$ \$/kW	
Lifespan:	$L^{13} = 10$	

Table S18: Electrolyzer parameters.

[17] At the very least it is plausible that pre-1990s progress in all technologies has played a roughly equivalent role in their overall progress in PtG systems. After all, they were all invented many decades ago, have had decades of experience at the R&D stage, and the main commercial application of 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.)

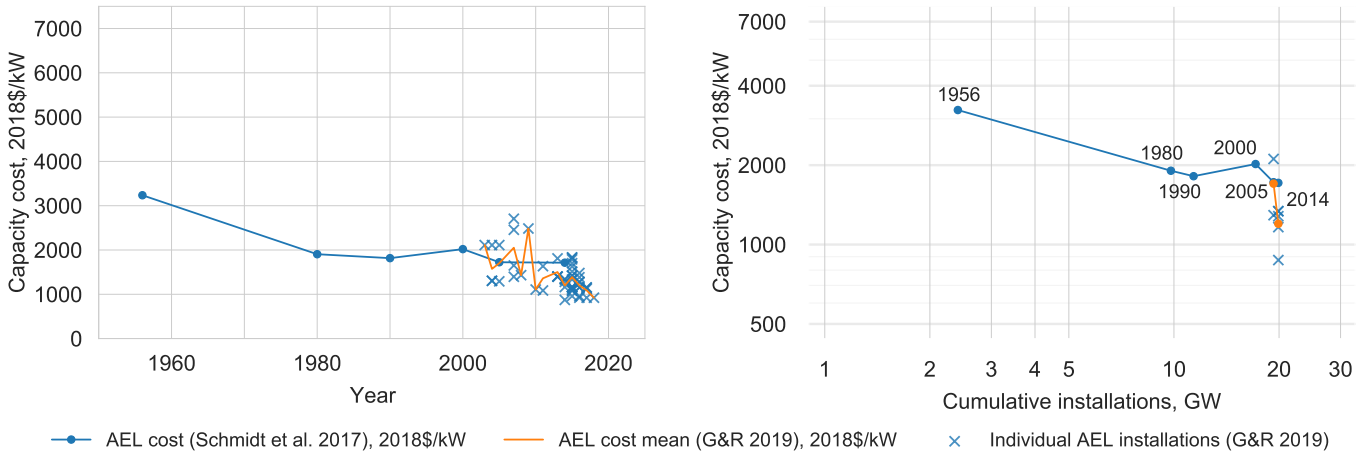

Figure S35: AEL electrolyzer cost time series (L) and experience curve (R).

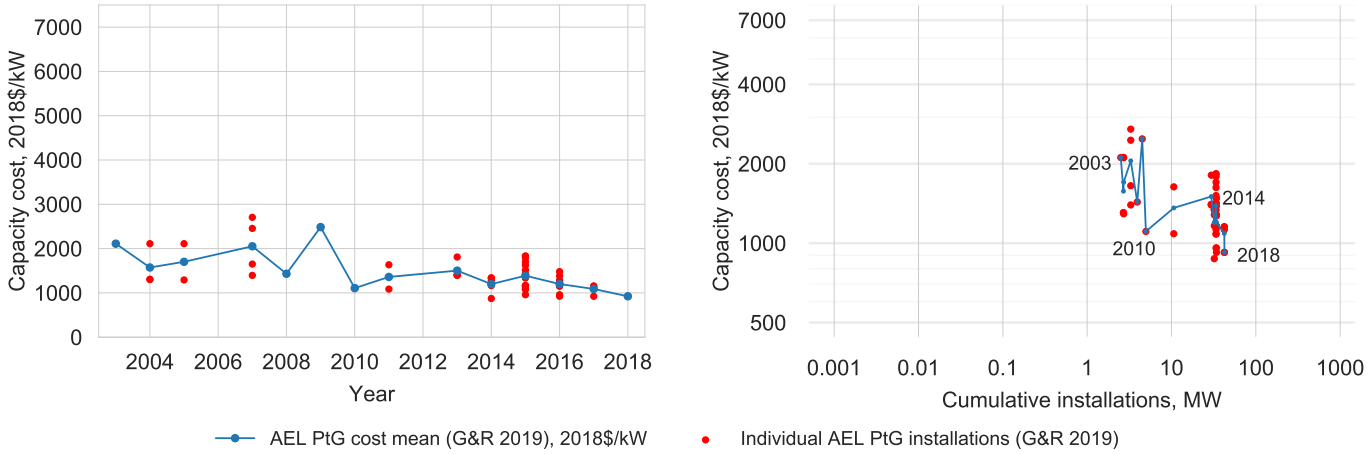

Figure S36: AEL PtG/PtL electrolyzer cost time series (L) and experience curve (R).

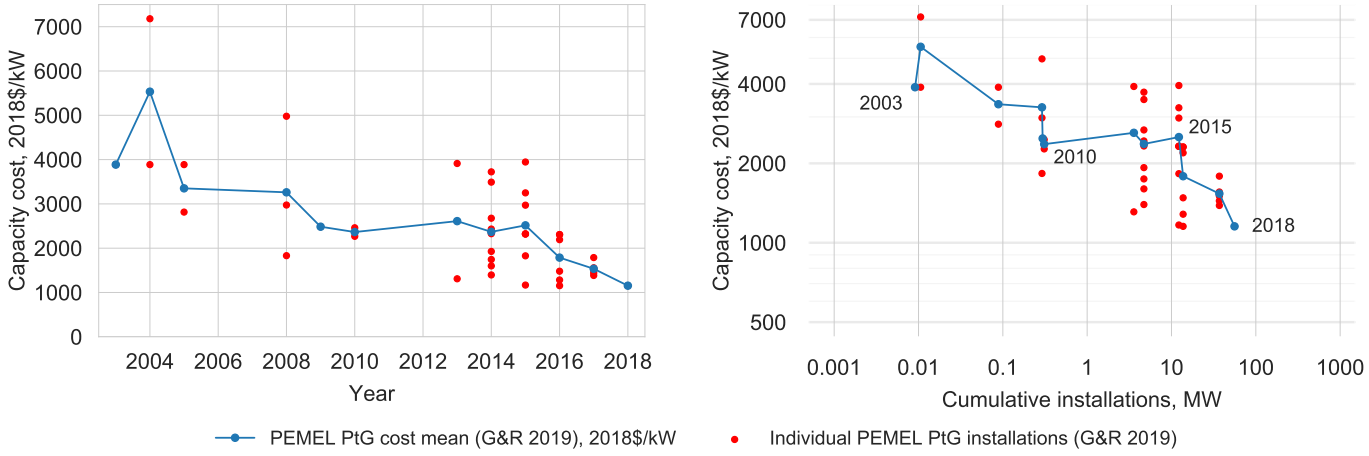

Figure S37: PEM PtG/PtL electrolyzer cost time series (L) and experience curve (R).

We observe that progress in PEM PtG has on average been faster than in AEL PtG, and the experience curve is steeper. 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 the mean 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 mean cost data

is slightly more successful, resulting in an estimated experience exponent of $\hat{\omega} = 0.194$ with standard error $\hat{\sigma}_\eta = 0.090$, and noise standard deviation $\hat{\sigma}_\eta = 0.233$, so although the uncertainty is still very 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 either encompasses SOEC technologies, or will lead to additional reversible functionality in future. 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 seems like a 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.

Now that we have chosen to use PEM/SOEC as our electrolyzer technologies, we must consider the initial capacity that exists and the initial hydrogen production per sector. The current installed capacity is around 140 MW for PEM in energy applications (and 0.3 MW for SOEC, which we ignore). This is approximately equal to the total cumulative production of these electrolyzers, because so few have been decommissioned. 140 MW is equivalent to 504 GJ/hour. At 50% capacity factor this is equal to around 0.002207 EJ/year. This is the electrical input though, so assuming a conversion efficiency of 60%[18] this gives around 0.0013245 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 evenly between the three end-use sectors to initialise the quantity of electrolytic hydrogen in each, from which to begin building out scenarios.

We take $\hat{\omega} = 0.194$, $\hat{\sigma}_\omega = 0.090$ and $\hat{\sigma}_\eta = 0.233$. Figure S38 shows production in each scenario with the associated cost forecasts given these parameters. We take the technology lifetime to be $L^{13} = 10$ years, and use an initial cost value of $c_0^{13} = 1100$ \$/kW. Parameter choices are summarized in Table S18.

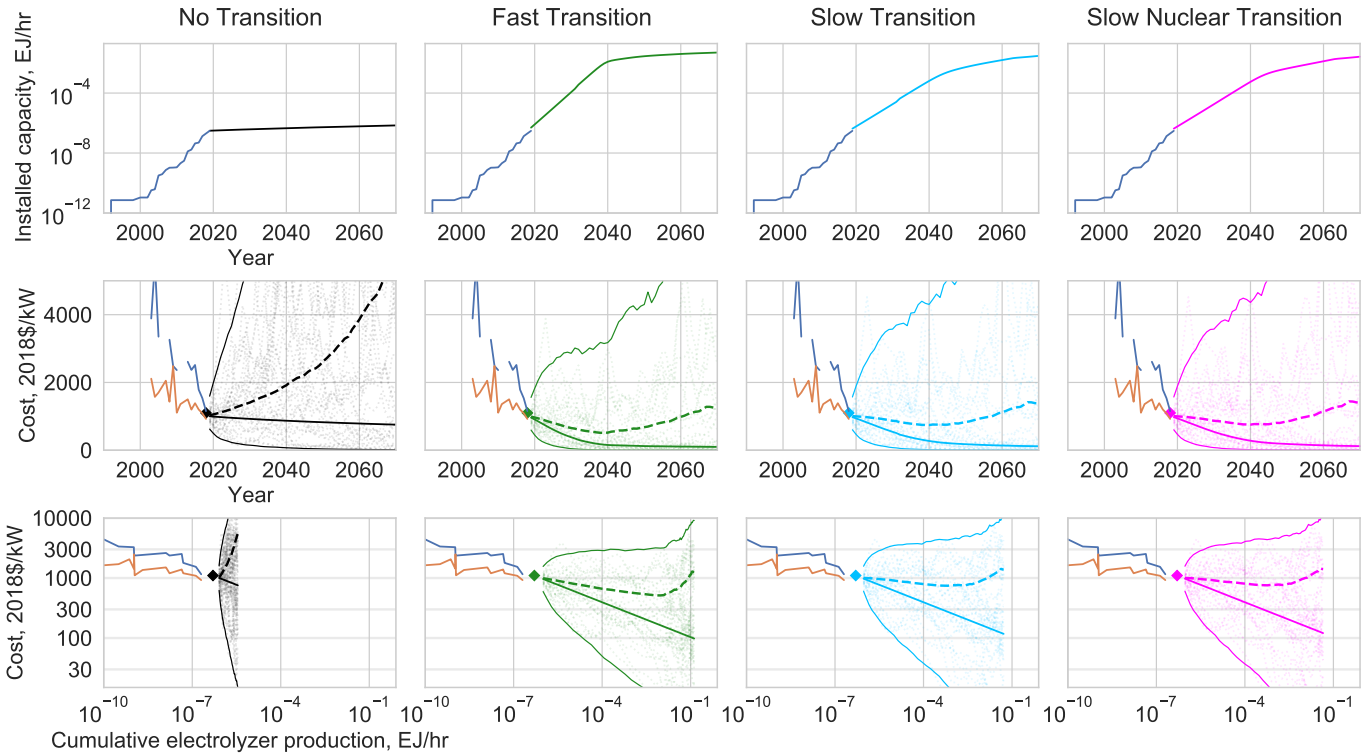

Figure S38: Electrolyzer installed capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figures S36 and S37.

[18] This conversion efficiency is what current PEM electrolyzers achieve. However, the long term efficiency of PEM is expected to be around 70%, which is the value used throughout the model, since it is only in the long term that electrolysis begins to play a large role. In addition, SOEC efficiency is currently around 80% and is expected to increase to around 90%, so there is plenty of opportunity for improvement.

6.13 Technology data summary

Table S19 summarises all the technology data used in the model, and Table S20 summarises the unit conversions.

Technology	Cost, c_0 (2018$)		Cumulative production, z_0		Annual production, E_0		Lifespan,	Cost model parameters			
	Std. units	$/GJ	Std. units	EJ	Std. units	EJ/yr	years				
AR(1) technologies								Mean-reversion $\hat{\phi}$	Constant $\hat{\kappa}$	Noise $\hat{\sigma}_\epsilon$	Long-term mean, $e^{\hat{\mu}}$
Oil[1]	61 $/bbl	10	-	-	3257 Mtoe	136.4	-	0.846	0.342	0.252	9.2
Coal[2]	56 $/short ton	2	-	-	920 Mtoe	38.5	-	0.954	0.035	0.090	2.1
Gas[3]	3 $/MMBtu	3	-	-	1346 Mtoe	56.3	-	0.824	0.207	0.237	3.3
Coal electricity	60 $/MWh	16.7	291180 TWh	1048	10101 TWh	36.4	-	0.927	0.206	0.102	16.4
Gas electricity	50 $/MWh	13.9	130320 TWh	469	6183 TWh	22.3	-	0.827	0.485	0.131	16.5
Wright's law technologies								Exp. exponent $\hat{\omega}$	Std. error $\hat{\sigma}_\omega$	Noise $\hat{\sigma}_\eta$	Learning rate, $1-2^{-\hat{\omega}}$
Nuclear	90 $/MWh	25	90280 TWh	325	2701 TWh	9.7	40	0.000	0.010	0.020	0
Hydropower	47 $/MWh	13	131944 TWh	475	4193 TWh	15.1	100	0.000	0.010	0.010	0
Biopower	72 $/MWh	20.0	9276 TWh	33.4	726 TWh	2.6	30	0.050	0.010	0.020	3.4%
Wind	55 $/MWh	15.3	8115 TWh	29.2	1270 TWh	4.6	30	0.158	0.045	0.103	10.4%
Solar PV	70 $/MWh	19.44	2220 TWh	8.0	584.6 TWh	2.1	30	0.303	0.047	0.093	18.9%
Daily batteries	360 $/kWh	100000	790 GWh	0.00284	-	-	10	0.342	0.073	0.116	21%
Multi-day storage	400 $/kWh	111100	0.1241 GWh	5×10^{-7}	-	-	20	0.168	0.045	0.103	11.0%
Electrolyzers	1100 $/kW	305555	140 MW	5×10^{-7} EJ/hr	-	-	10	0.194	0.090	0.233	12.6%

Table S19: Technology parameters used in main case analysis.

[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	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 \$/boe = $\frac{1}{6.12}$ \$/GJ = 0.163 \$/GJ Average price in 2019 (BP 2019): 71.31 \$/bbl This is equivalent to approximately 11.7 \$/GJ = 11.7 bn\$/EJ
Oil refining	EIA (2019*c*): average retail price of gasoline is: 52% crude oil, 11% refining costs and profits 19% taxes, 18% distribution and marketing $\therefore$ at 71.31 \$/bbl, refining costs and profits are $\frac{11}{52}\times$ 71.31 \$/bbl $\approx$ 15 \$/bbl Assume 50-50 split between refining costs and profits $\Rightarrow$ Refining costs approximately 7.5 \$/bbl, or $0.5\times\frac{11}{52}\times$ 11.7 \$/GJ $\approx$ 1 bn\$/EJ (Estimate is conservative due to overestimating refining profits and rounding down.)
Bituminous coal (i.e. non-power-sector coal)	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 (EIA 2019*a*): 59.43 \$/short ton = 65.51 \$/metric ton Anthracite price in 2018, US (EIA 2019*a*): 99.97 \$/short ton = 110.2 \$/metric ton Price in 2018, Australia (FRED 2019): 113.23 \$/metric ton Bituminous US value is equivalent to 2.18 \$/GJ
Natural gas	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
Electricity	1 TWh = 0.0036 EJ 1 MWh = 3.6 GJ 1 kWh = 3.6 MJ 277.8 TWh = 1 EJ $\therefore$ 1 \$/MWh = $\frac{1}{3.6}$ \$/GJ = 0.2778 \$/GJ and 1 \$/kWh = $\frac{1}{3.6}$ \$/MJ = 0.2778 \$/MJ = 277.8 \$/GJ E.g. LCOE: 50 \$/MWh = 13.90 \$/GJ = 13.90 bn\$/EJ E.g. EV battery packs: 176 \$/kWh $\approx$ 49000 \$/GJ E.g. Li-ion grid storage: 360 \$/kWh $\approx$ 100000 \$/GJ E.g. VRF grid storage: 400 \$/kWh $\approx$ 111000 \$/GJ
Electrolyzers	Current cost is around 1100 \$/kW 1 kW = 1 kWh/h = 10^{-9} TWh/h = 3.6×10^{-12} EJ/h $\therefore$ 1100 \$/kW = $1100\times10^{12}/3.6$ \$/(EJ/h) $\approx$ 305555 \$/(GJ/h)

Table S20: 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 $T = 52$ the resulting ensemble consists of 692 values.

We generated this complete ensemble $M = 100000$ times, with the parameters shown in Table S19 (this takes around 50 minutes on the 2015 MacBook Pro used). For each ensemble we computed the implied cost sample path for each technology (shown in Section 6, e.g. Figures S14 and S22). We then calculated the annual system cost sample at each time for each scenario (Eq. 38), thus obtaining the annual cost distributions, $W_t^{scenario}$. Figure S39 shows the evolution of these distributions for the four scenarios, and Figure S40 shows the median annual spend on each technology in each scenario.

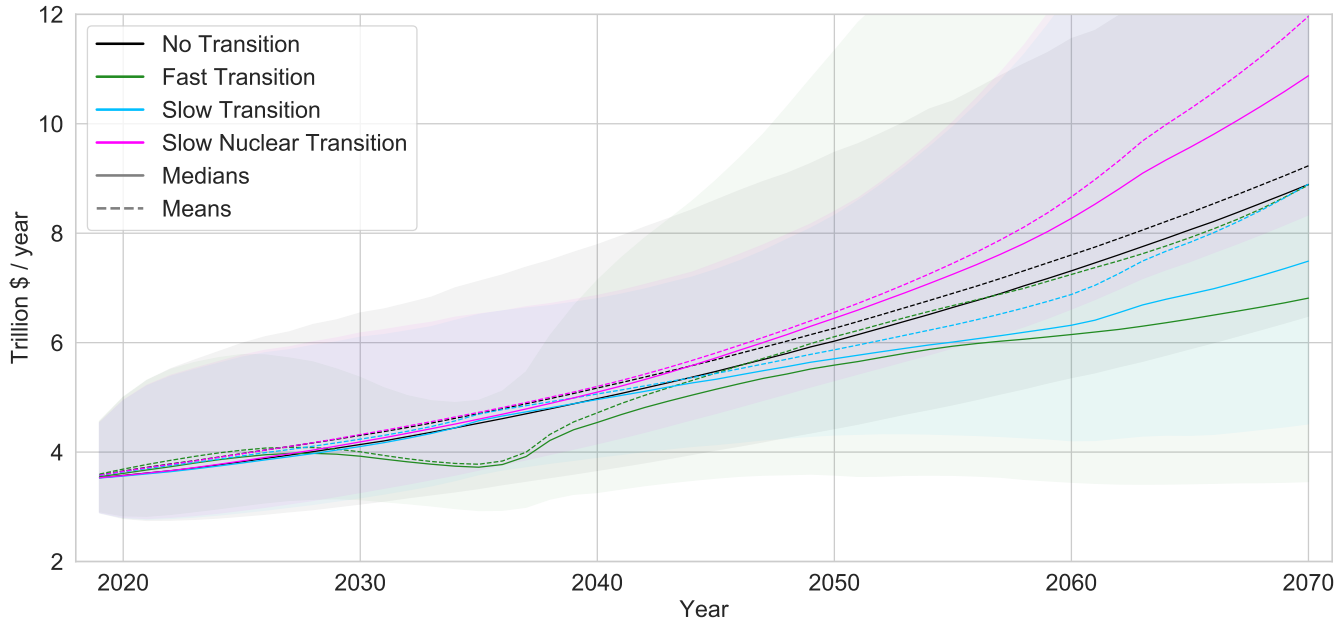

Figure S39: Annual system cost distributions, $W_t^{scenario}$, for each scenario. Medians are shown as solid lines, expectations as dashed lines, and the shaded regions represent the 95% confidence intervals.

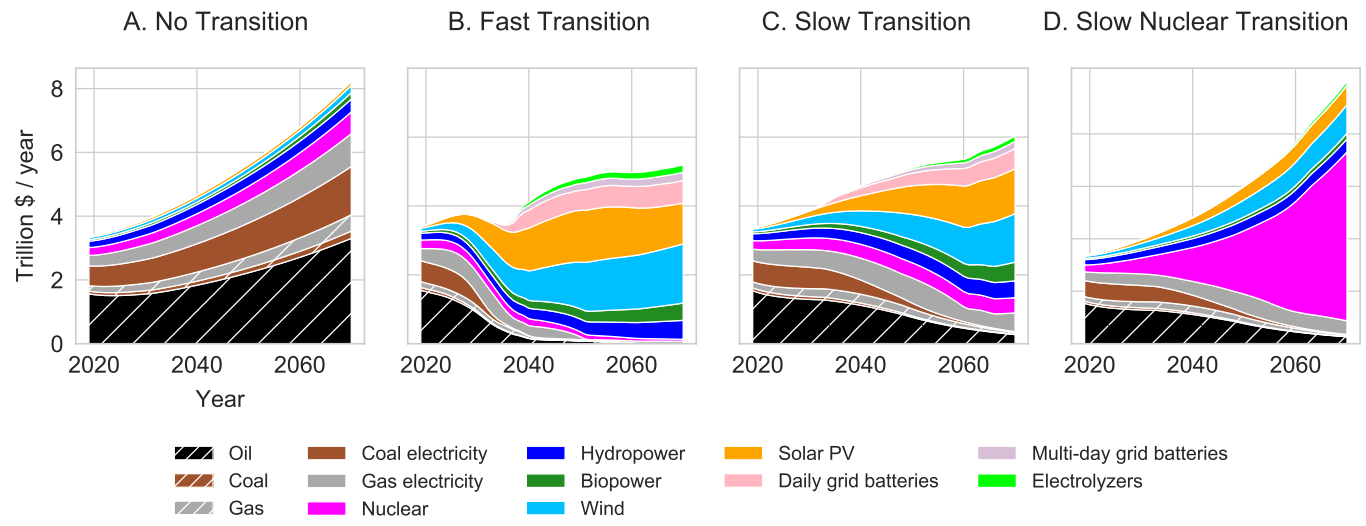

Figure S40: Median annual spend on each technology in each scenario.

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-decarbonize 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 Section 5.2), these diverge strongly.

We then picked a discount rate r and, for each of the M ensembles, computed the implied present discounted system cost sample for each scenario (Eq. 39), 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 S41 (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 NPC distributions 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 NPC of each scenario varies with r. This is shown in Figure S41 (R). We also observe how the median NPC of each scenario varies with r, as shown in Figure S42 (R).

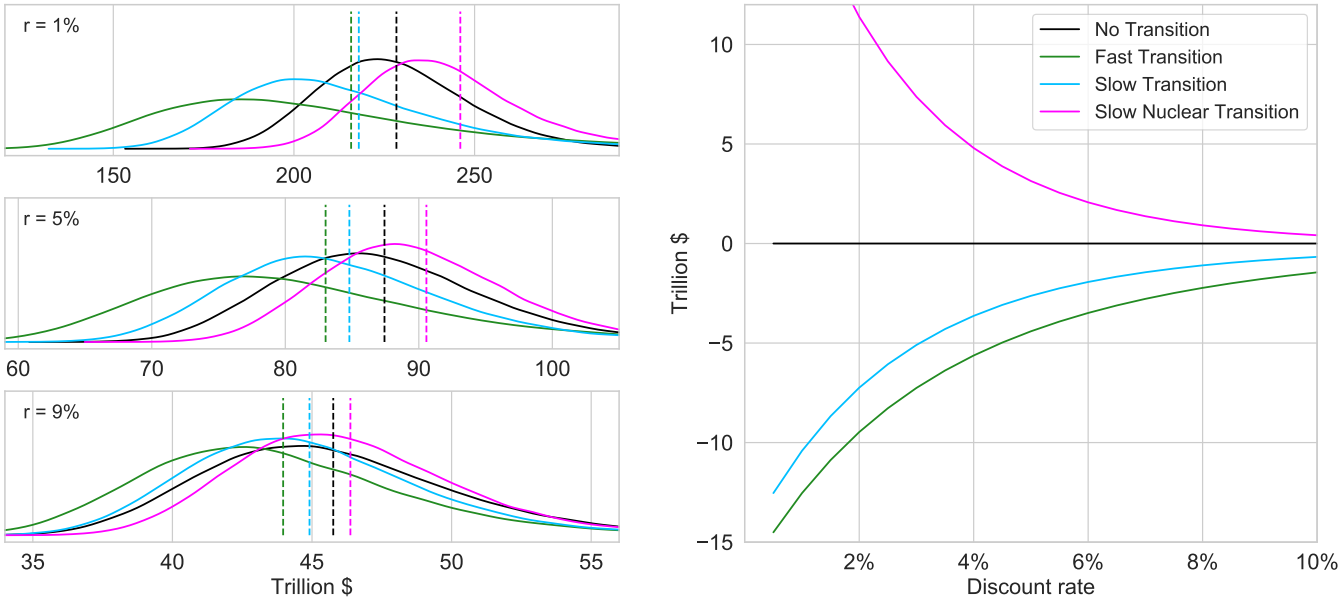

Figure S41: Present discounted cost distribution of each scenario at varying discount rates, with expected values shown as dashed lines (L). Expected net present cost of transition for each scenario relative to No Transition (R).

The expected net present cost of a rapid green energy transition is negative at all discount rates up to around 25% Therefore, in terms of energy system and technology costs alone, pursuing this course of action is a bet worth taking. Furthermore, this analysis include no climate damages, so the full economic costs of No Transition, Slow Transition and Slow Nuclear Transition 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).

7.2 Dependence on 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 (with $M = 20000$ to save computation time) are shown in Figure S43.

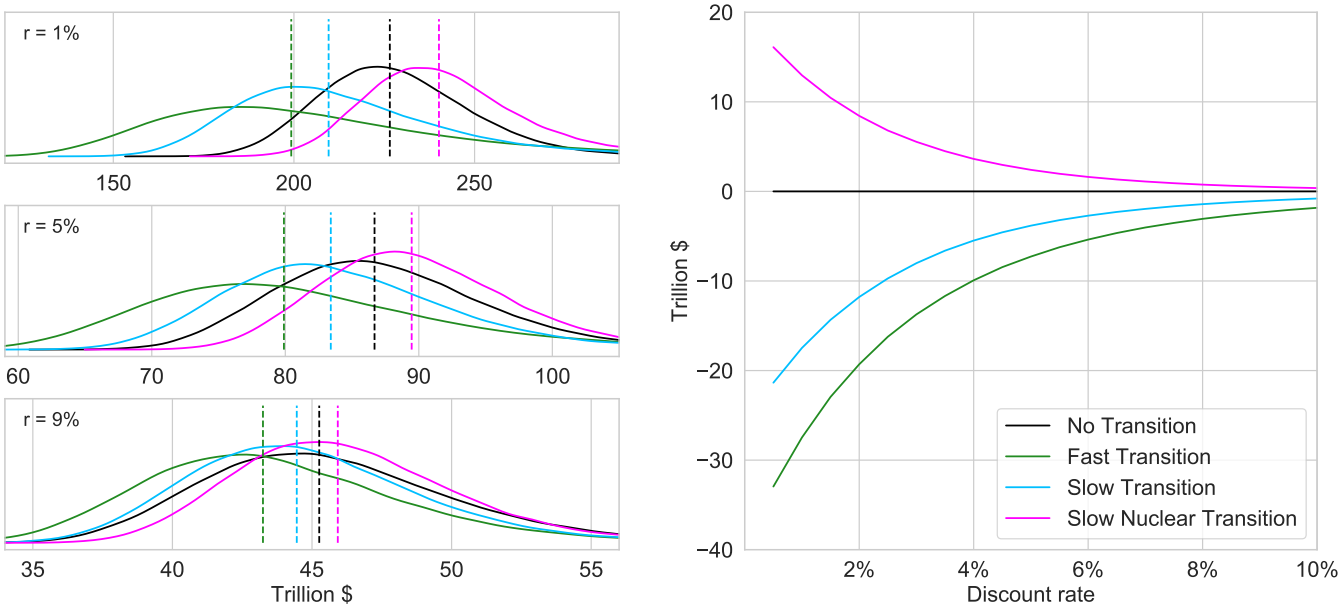

Figure S42: Present discounted cost distribution of each scenario at varying discount rates, with median values shown as dashed lines (L). Median net present cost of transition for each scenario relative to No Transition (R).

7.3 Side cases

To explore how results depend on technology parameters, we present a few side cases reflecting alternative parameter choices: 1) Less pessimistic wind and solar costs; 2) Lower oil and gas prices; 3) Lower nuclear costs; 4) Less uncertain electrolyzer costs; and 5) a combination of 1, 2 and 4. In each case only the stated parameters are changed, all others remain identical to the main case. In the following analyses we took $M = 20000$ again, to reduce computation time.

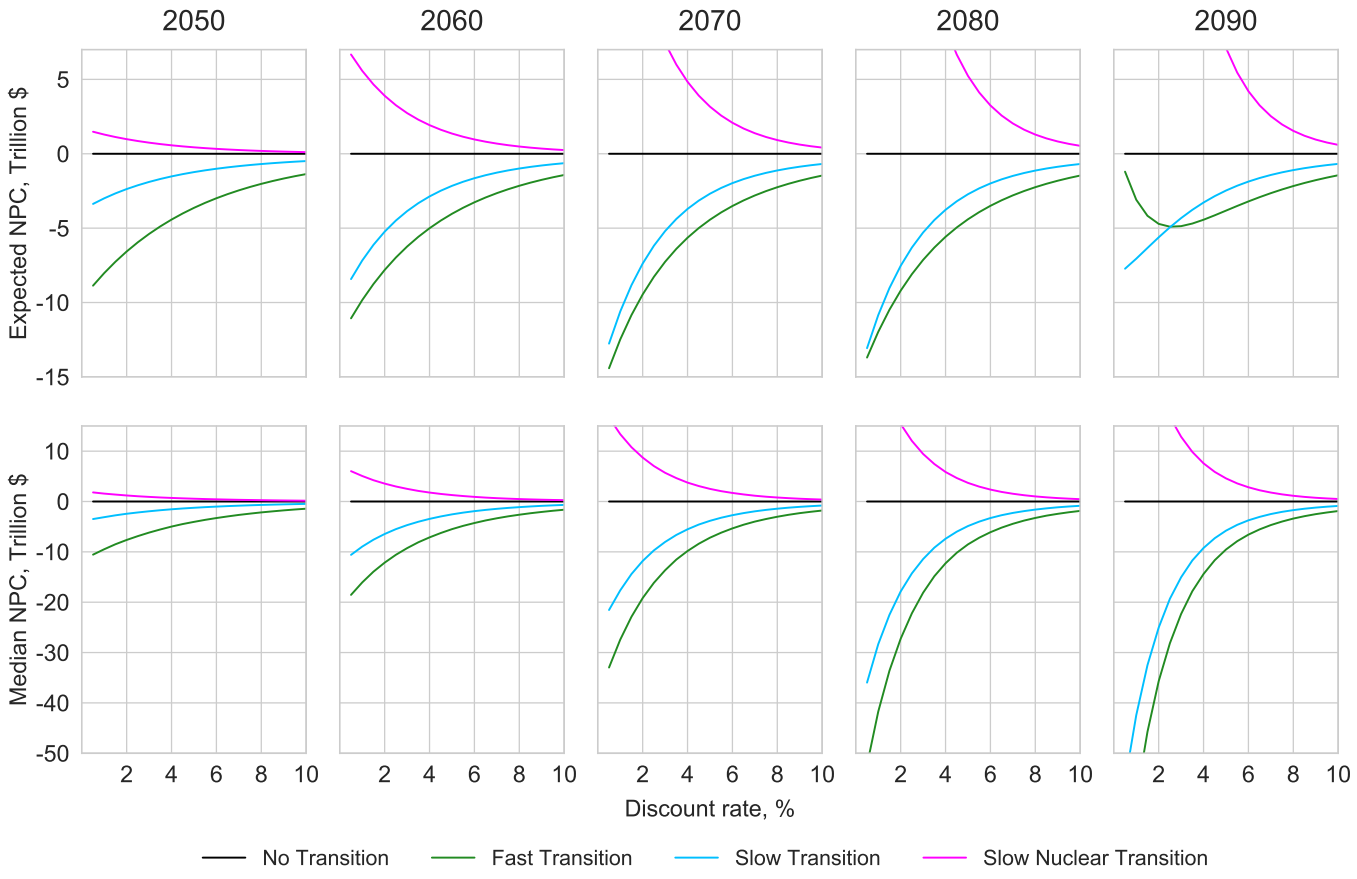

Figure S43: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090.

7.3.1 Less pessimistic solar and wind costs

The main case was initialised with a current cost of solar electricity of $c_0^{10} = 70$ \$/MWh, which is higher than many current estimates. We used this value because although most current and recent solar LCOE estimates are around 70 \$/MWh or lower (many around 40 \$/MWh), the most recent number from IRENA is 85 \$/MWh. In this side case we use the value $c_0^{10} = 50$ \$/MWh. For wind the difference isn't so pronounced; the main case again used $c_0^9 = 55$ \$/MWh, and in this side case we use $c_0^9 = 50$ \$/MWh. BNEF's most recent 2019 global benchmark values are 47 \$/MWh and 51 \$/MWh, so this side case seems entirely plausible.

Wind and 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), 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 huge effect. The cost forecasts for wind and solar in this side case are shown in Figures S44 and S45, and the NPC results are shown in Figure S46. At 4% discount rate the Fast Transition scenario consistently results in a NPC saving of around 8-12 trillion \$.

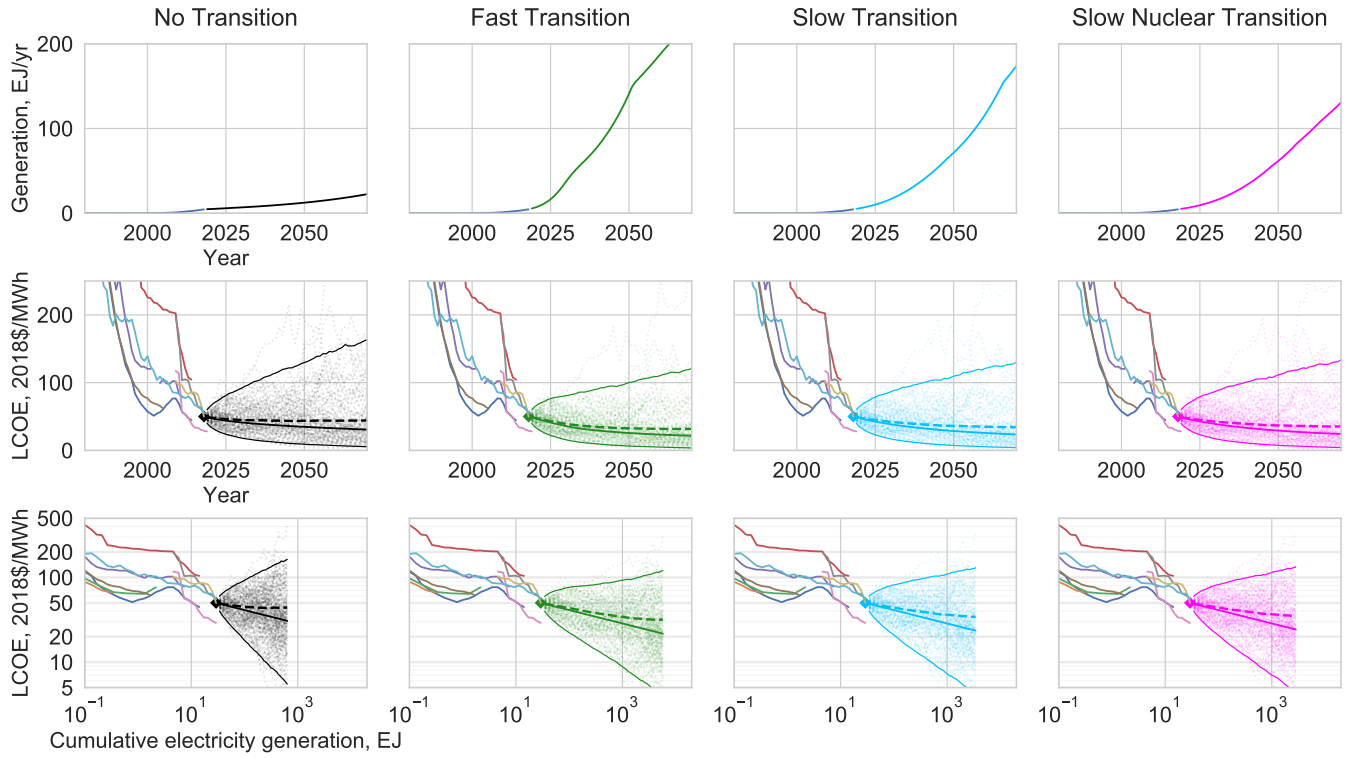

Figure S44: Wind electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row) when initialised with less pessimistic current cost $c_0^9 = 50$ \$/MWh (instead of 55 \$/MWh in main case). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S28.

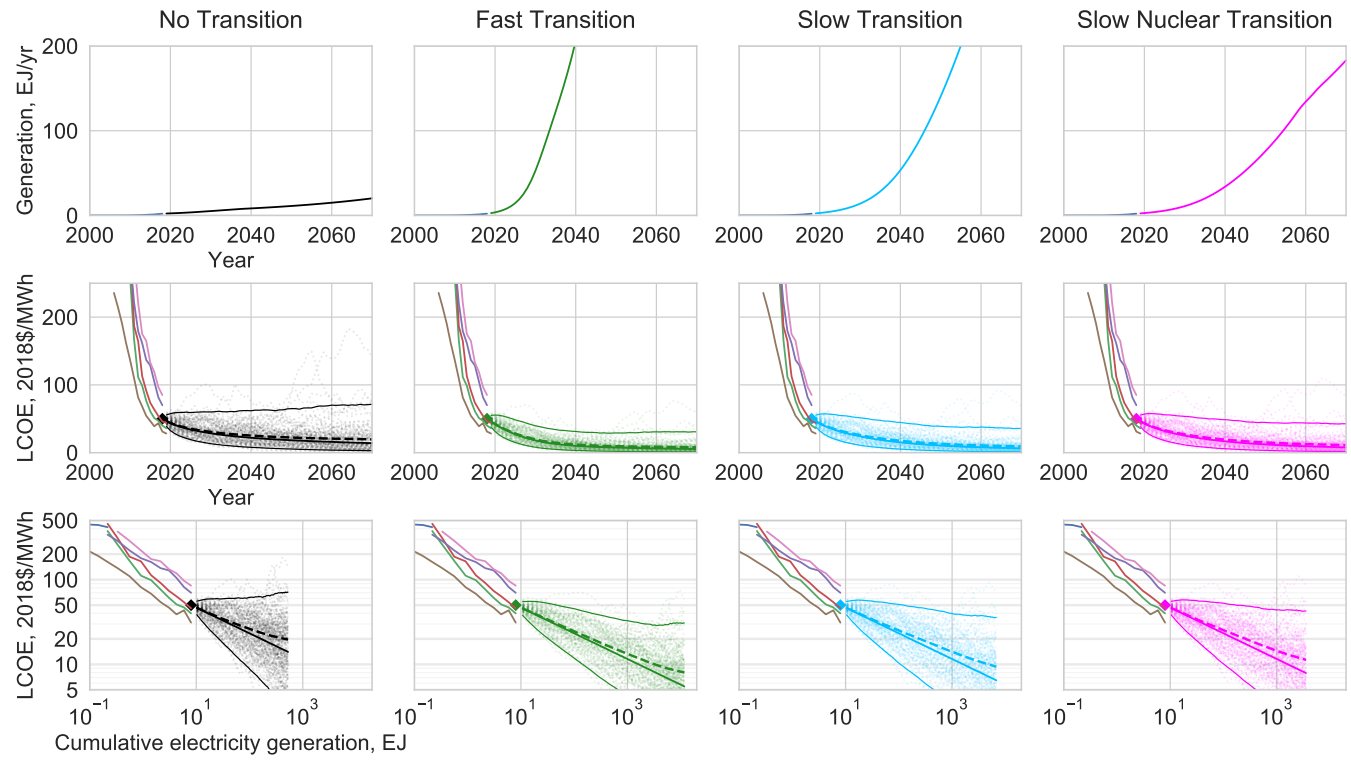

Figure S45: Solar PV electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row), when initialised with less pessimistic current cost c_0^{10} = 50 $/MWh (instead of 70 $/MWh in main case). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S30.

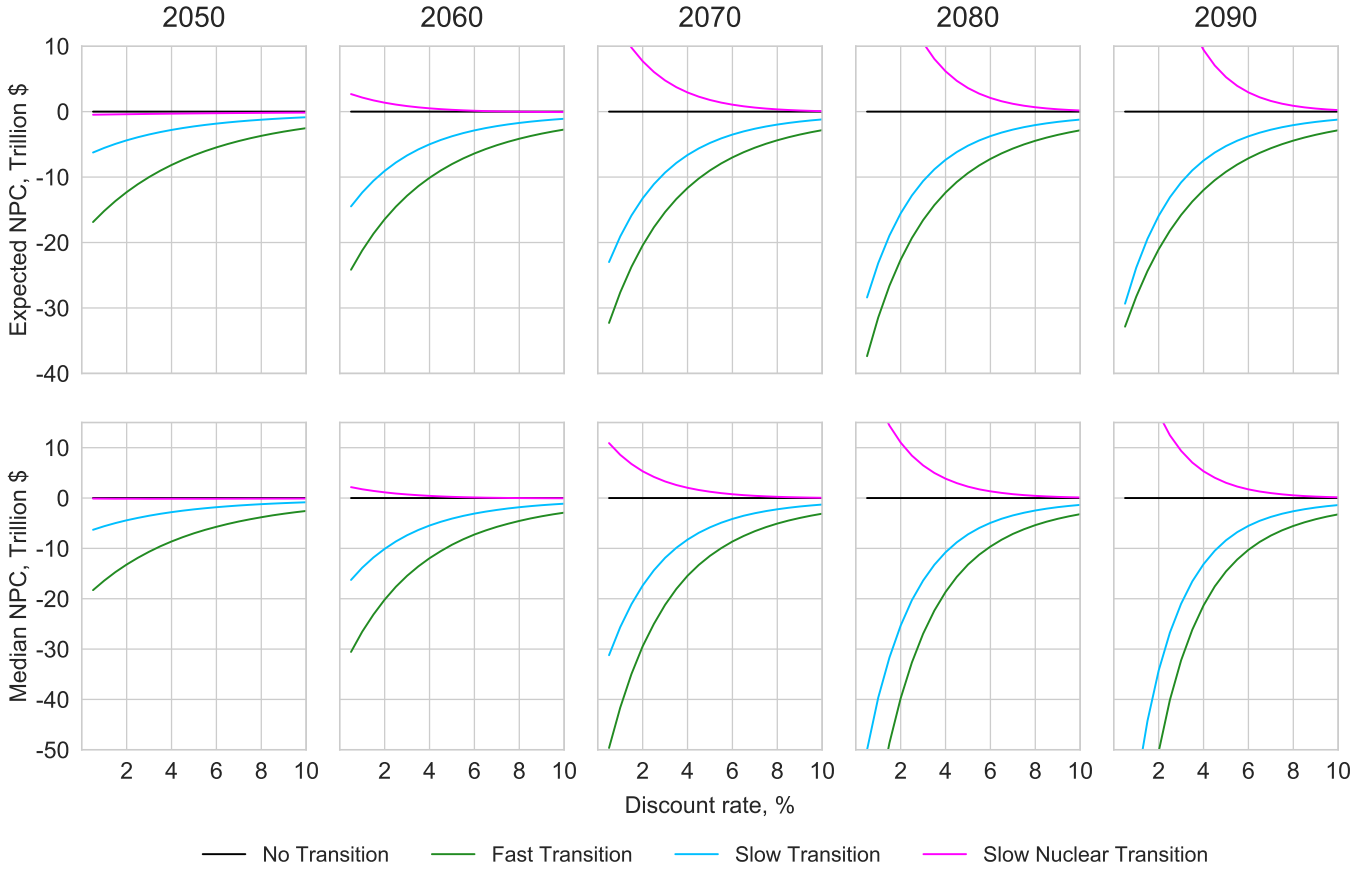

Figure S46: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with less pessimistic wind and solar current costs

7.3.2 Lower oil and gas prices

Recall, the AR(1) models for oil and gas were calibrated on price data from the last four decades, from 1975 onwards. If instead we use the entire historical data sets, back to 1861 and 1922 respectively, the estimated parameters are for oil: $\widetilde{\phi} = 0.908$, $\widetilde{\sigma}_\epsilon = 0.266$ and $\widetilde{\kappa} = 0.145$; and for gas: $\widetilde{\phi} = 0.973$, $\widetilde{\sigma}_\epsilon = 0.177$ and $\widetilde{\kappa} = 0.014$. 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 S47 and S48. 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 S49. In this case all transition paths have higher expected NPC than remaining with the current system, though with a long enough time horizon the median NPC values eventually fall below those of No Transition.

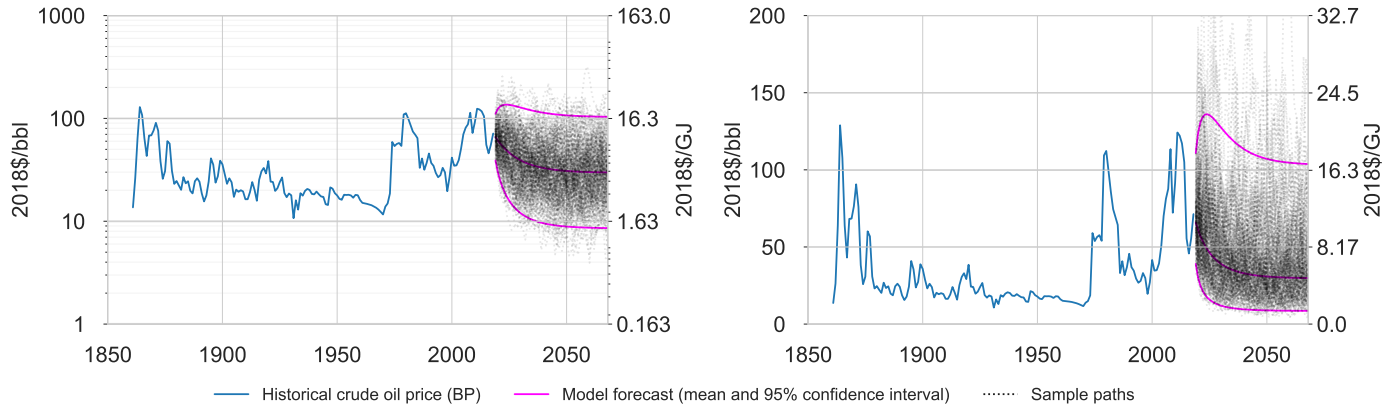

Figure S47: Low oil price forecast resulting from calibration on data 1861-2018, analogous to Fig. S13.

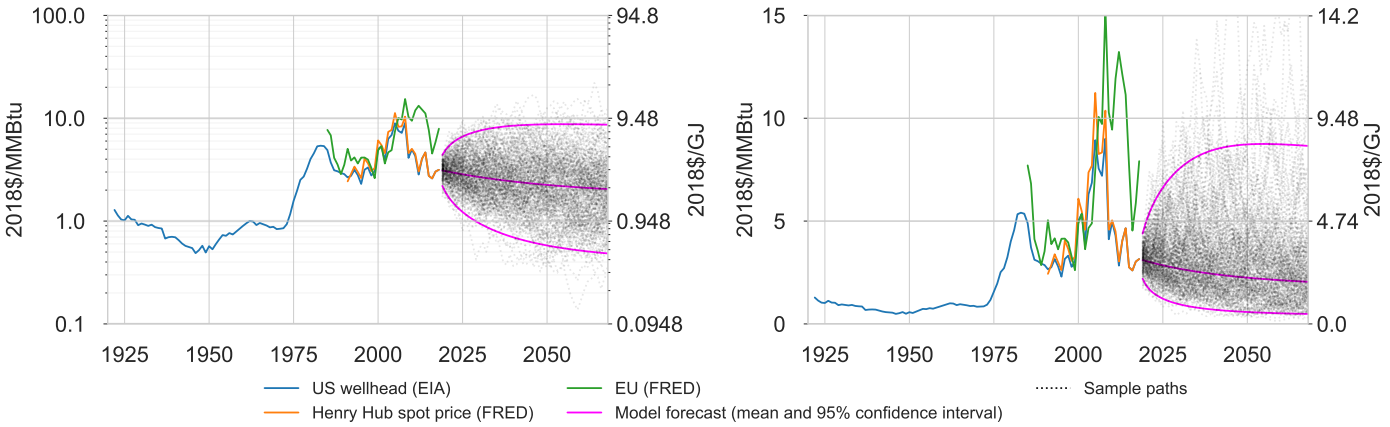

Figure S48: Low gas price forecast resulting from calibration on data 1922-2018, analogous to Fig. S15.

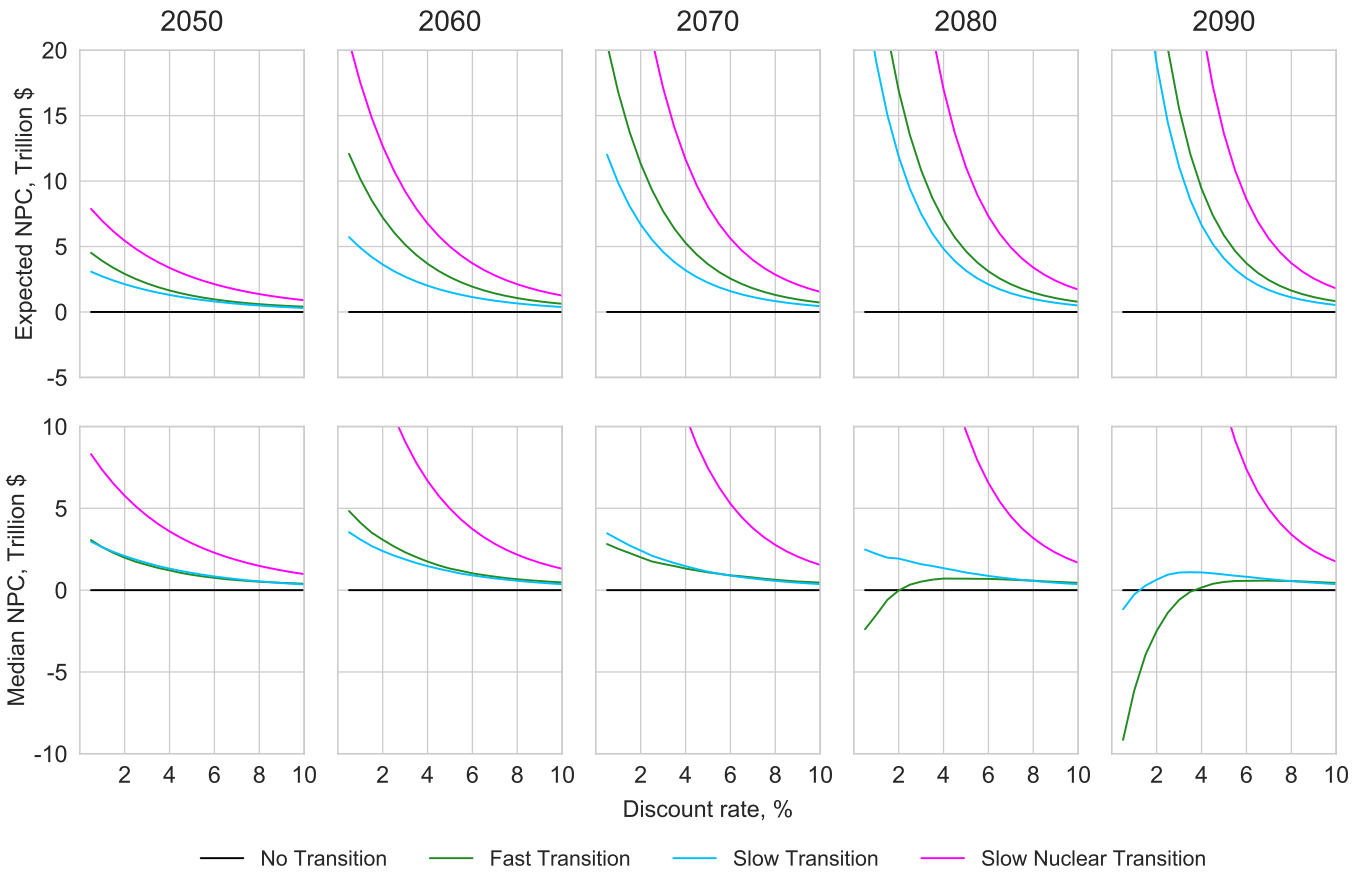

Figure S49: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with low oil and gas prices.

7.3.3 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 an arbitrary assumption. The cost forecasts for each scenario are shown in Figure S50 and the NPC results are shown in Figure S51. Total spend on nuclear is now much lower, making the Slow Nuclear Transition more competitive generally, but it still does not beat Fast Transition.

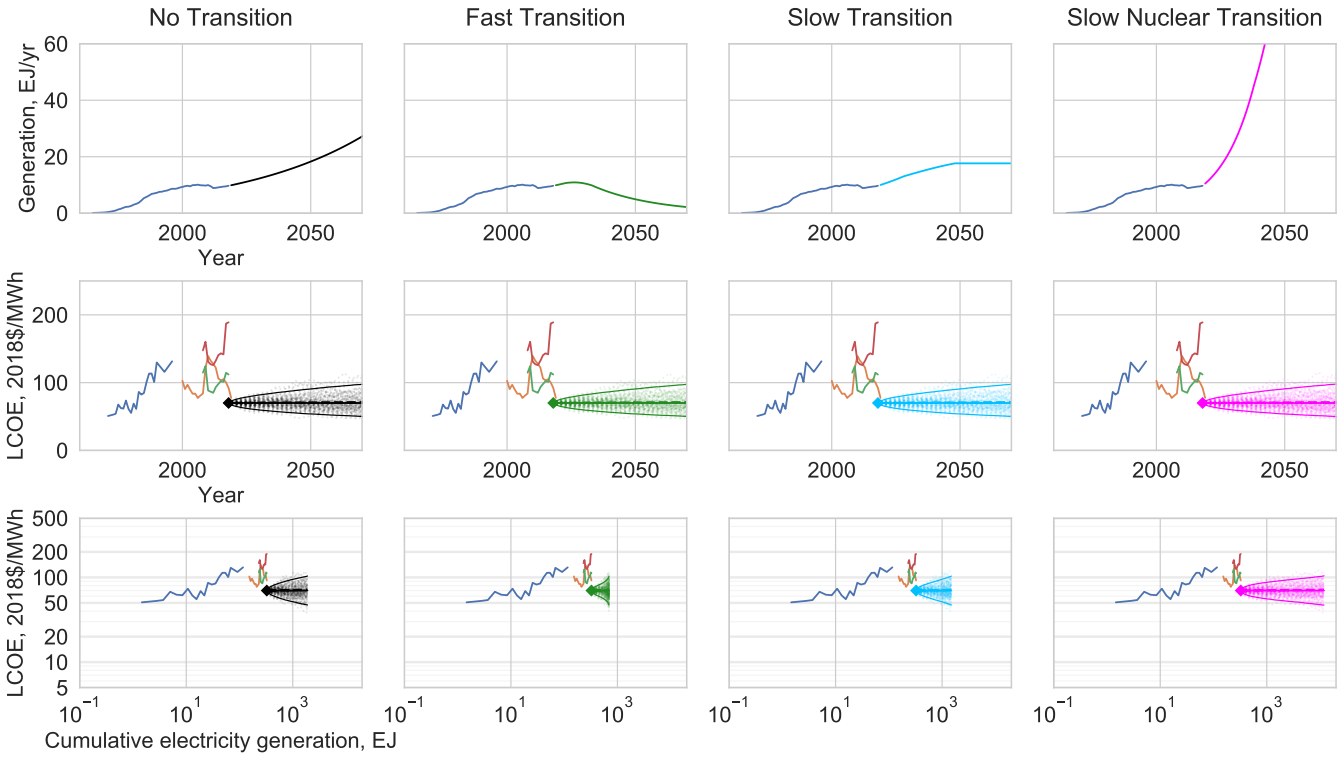

Figure S50: Nuclear electricity generation in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row), when initialised with low current cost $c_0^6 = 70$ \$/MWh (instead of 90 \$/MWh in main case). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Figure S20.

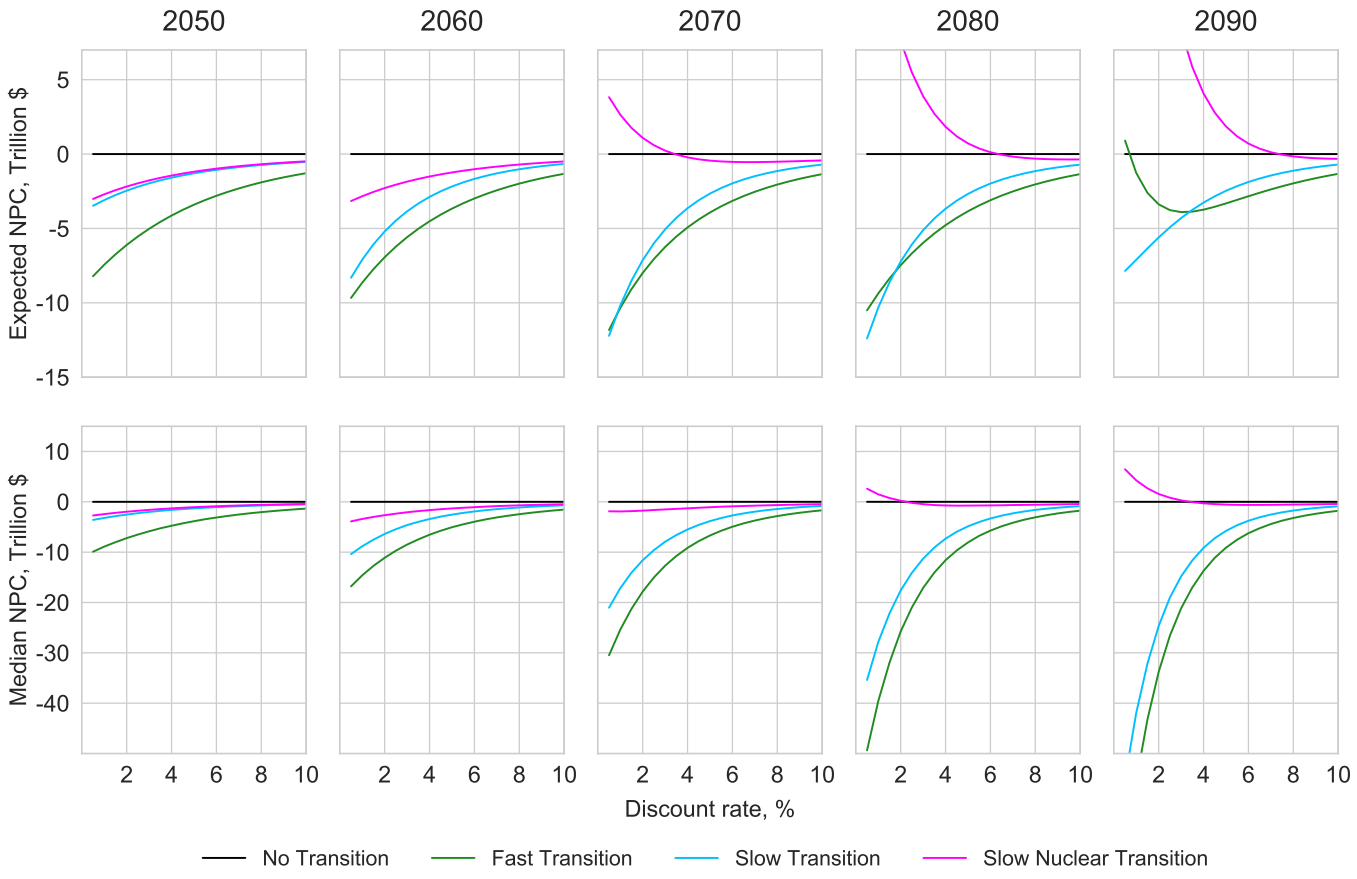

Figure S51: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with low nuclear costs.

7.3.4 Less uncertain electrolyzer costs

Based on calibrating the FDWL model on historical data, the main case uses parameters $\hat{\omega} = 0.194$, $\hat{\sigma}_\omega = 0.090$ and $\hat{\sigma}_\eta = 0.233$. These parameters give forecasts with very strong median progress but very large uncertainty. This is a direct result of the small sample size of the historical data, as this causes the series of mean values to have large fluctuations. This side case considers the situation where the progress trend is much weaker but the uncertainty is also much lower. We halve the experience exponent to $\hat{\omega} = 0.100$, halve its standard error to $\hat{\sigma}_\omega = 0.040$, and reduce the noise standard deviation to $\hat{\sigma}_\eta = 0.080$. These parameters are more like those observed in other technologies generally. We keep the current initial cost of 1100 $/kW. The forecasts for each scenario when these parameters are used are shown in Figure S52. Note that in all but No Transition, costs around 2040 are in the range 100-1500 $/kW, with medians around 300-800 $/kW. As noted earlier, current estimate 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 S53. 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 2070.

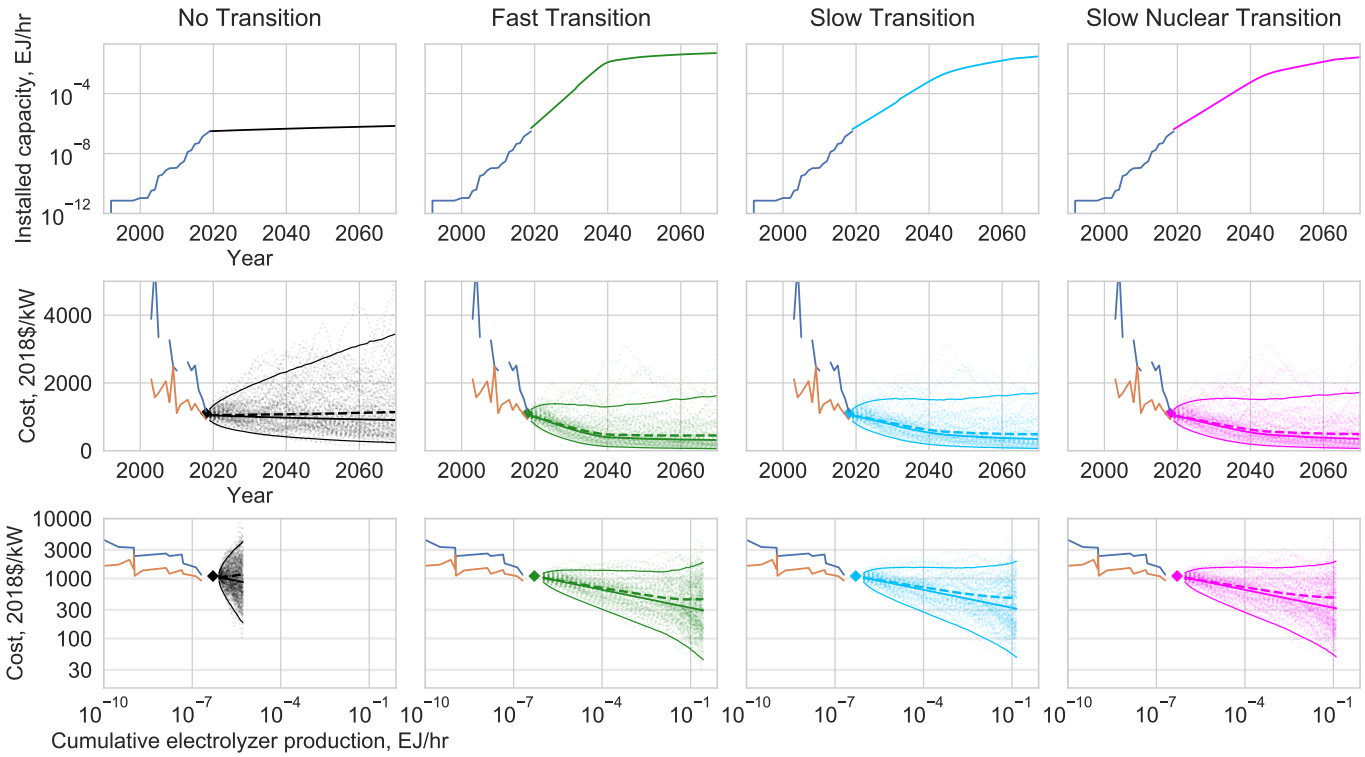

Figure S52: Electrolyzer installed capacity in each scenario (top row), with the corresponding cost forecast distributions plotted against time (middle row) and experience (bottom row), when initialised with slower progress but less uncertainty: $\hat{\omega} = 0.100$, $\hat{\sigma}_\omega = 0.040$, $\hat{\sigma}_\eta = 0.080$ (instead of $\hat{\omega} = 0.194$, $\hat{\sigma}_\omega = 0.090$, $\hat{\sigma}_\eta = 0.233$ 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 S36 and S37.

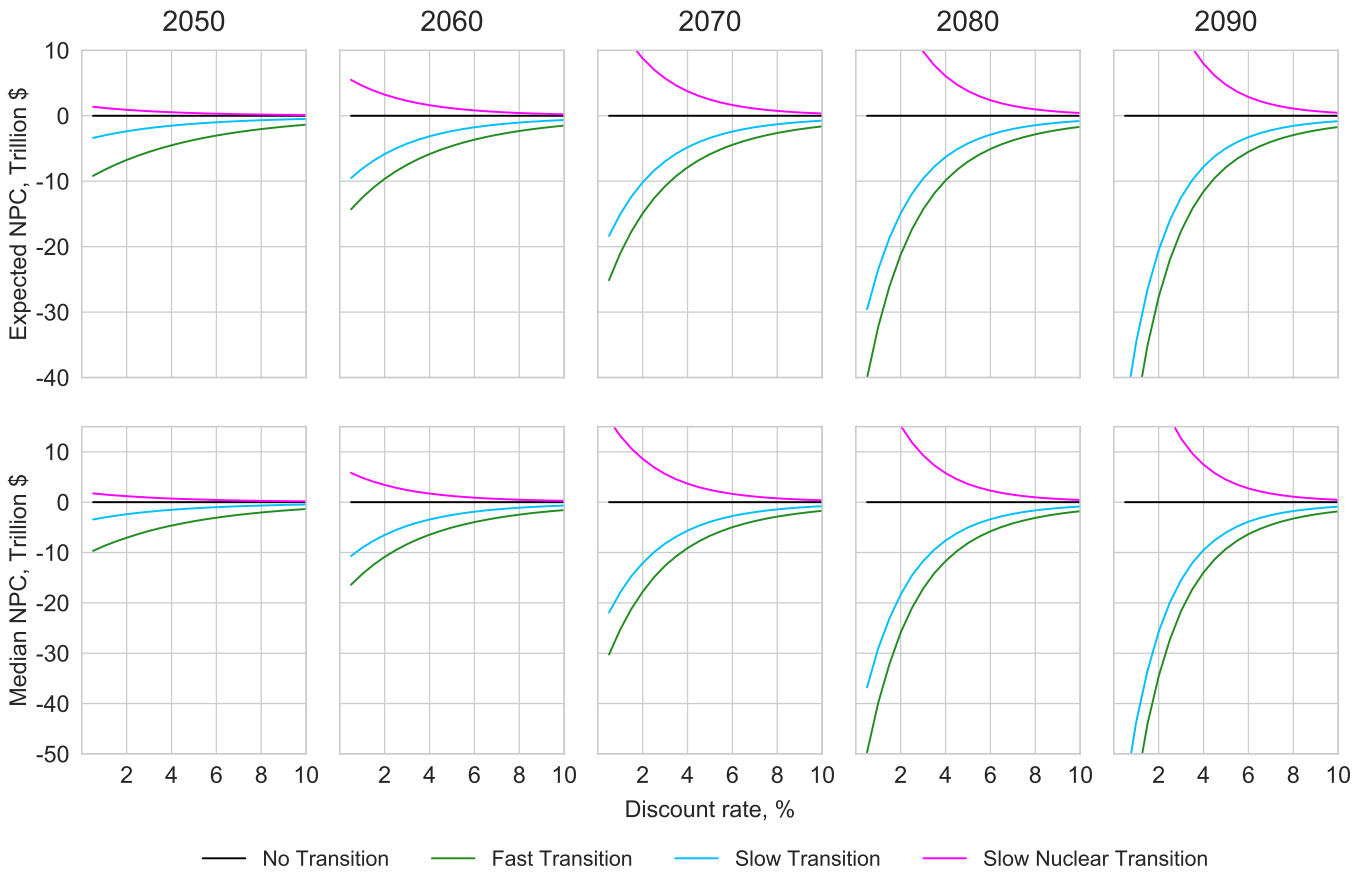

Figure S53: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with less uncertain electrolyzer costs.

7.3.5 Less pessimistic wind and solar costs, low oil and gas prices and less uncertain electrolyzer costs

This side case is the combination of three of the above side cases. We suppose that *i*) oil and gas prices return to low, pre-globalisation levels, *ii*) the most up to date global average wind and solar cost estimates made by industry analysts are accurate, and *iii*) electrolyzer costs roughly follow the path expected by manufactures, as the industry scales up, and do not instead escalate dramatically. The NPC results are shown in Figure S54. In this case, we find that the Fast Transition scenario has negative mean and median NPC at all discount rates and time horizons investigated.

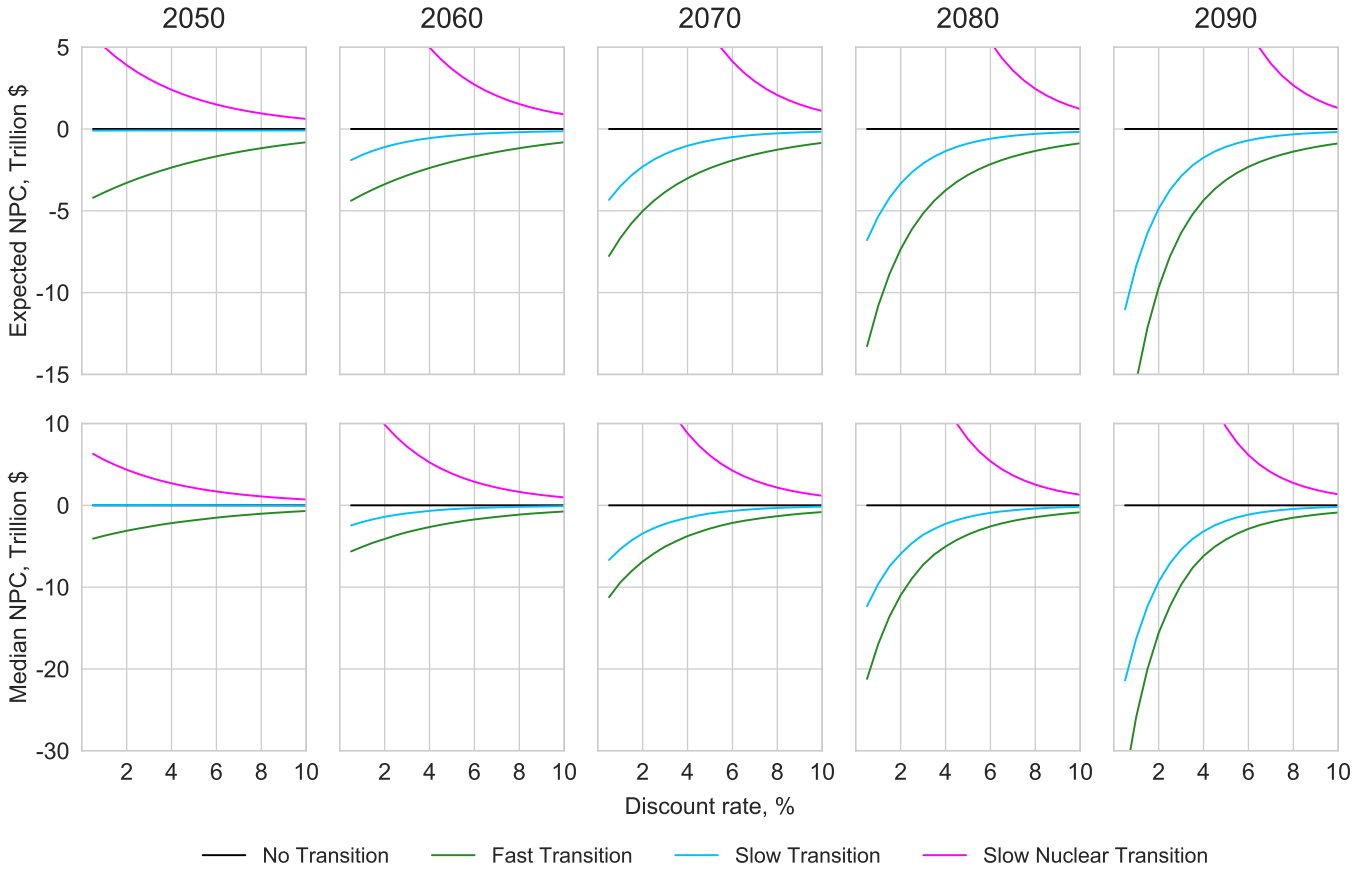

Figure S54: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2050 and 2090, with less pessimistic wind and solar costs, low oil and gas costs, and less uncertain electrolyzer costs.

7.4 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 analyzed in Farmer & Lafond (2016). The analysis so far has used the first-difference Wright's law model given in Eq. 31; the analogous first-difference Moore's law (FDML) equation is

$$\log c_t - \log c_{t-1} = -\mu + n_t \qquad \text{with IID } n_t \sim \mathcal{N}(0, \sigma_n^2), \tag{43}$$

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,

$$\log c_t - \log c_{t-1} = -\mu + v_t + \theta v_{t-1} \qquad \text{with IID } v_t \sim \mathcal{N}(0, \sigma_v^2), \tag{44}$$

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

$$\sigma_v^2 = \sigma_n^2/(1+\theta^2). \tag{45}$$

Farmer & Lafond (2016) 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. 43) on historical data we obtain estimates for the three values $\hat{\mu}$, $\hat{\sigma}_\mu$ and $\hat{\sigma}_n$. Eq. 45 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. 44 to compute c_t at each step. Repeating this sampling process a large number of times recovers the theoretical forecast distribution calculated in Farmer & Lafond (2016). 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 Sections 6.6 - 6.12, 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 S55. This model is not well suited to forecasting beyond about 30 years for these tehcnologies, as the forecast variance becomes very large. However, since the difference in production of these three technologies over the four 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 Moore's law parameters for the different wind series presented in Section 6.9 are shown in Table S21. We use the values from the chained series at the bottom of the table: $\hat{\mu} = 0.083$, $\hat{\sigma}_\mu = 0.015$, $\hat{\sigma}_n = 0.086$.

Data source	$\hat{\mu}$	$\hat{\sigma}_\mu$	$\hat{\sigma}_n$	Start year	Final year	Series length	Current LCOE
Original data							
US DOE (Wiser et al. 2016)	0.077	0.013	0.077	1980	2014	35	
BNEF 2015 (Wiser et al. 2016)	0.069	0.013	0.074	1984	2014	31	
BNEF 2011 (Wiser et al. 2016)	0.049	0.018	0.089	1984	2011	26	
Schilling & Esmundo (2009) (mean)	0.091	0.008	0.039	1980	2005	26	
Lazard LCOE Analysis Low	0.143	0.065	0.206	2009	2019	11	28 $/MWh
Lazard LCOE Analysis High	0.129	0.047	0.149	2009	2018	11	54 $/MWh
Bloomberg Terminal / BNEF	0.080	0.037	0.111	2009	2018	10	50 $/MWh
IRENA (2018*b*)	0.059	0.012	0.073	1983	2017	35	56.6 $/MWh
IRENA (2019)	0.049	0.010	0.061	1983	2018	36	55 $/MWh
Chained data							
BNEF 2015 + Bloomberg Term.	0.083	0.015	0.086	1984	2018	35	50 $/MWh

Table S21: Moore's law parameters for wind (*c.f.* Table S13).

The Moore's law parameters for the different solar PV series presented in Section 6.10 are shown in Table S22. We use the values from the chained series at the bottom of the table: $\hat{\mu} = 0.124$, $\hat{\sigma}_\mu = 0.017$, $\hat{\sigma}_n = 0.109$.

The FDML forecasts for wind and solar are shown in Figure S56. 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. Again it appears that FDML is not appropriate for forecasting beyond about 30 years.

For Li-ion batteries we use the long term trend labelled "Li-ion (all applications)" on Figure S32 for calibration, yielding parameter estimates: $\hat{\mu} = 0.128$, $\hat{\sigma}_\mu = 0.021$, $\hat{\sigma}_n = 0.099$.

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

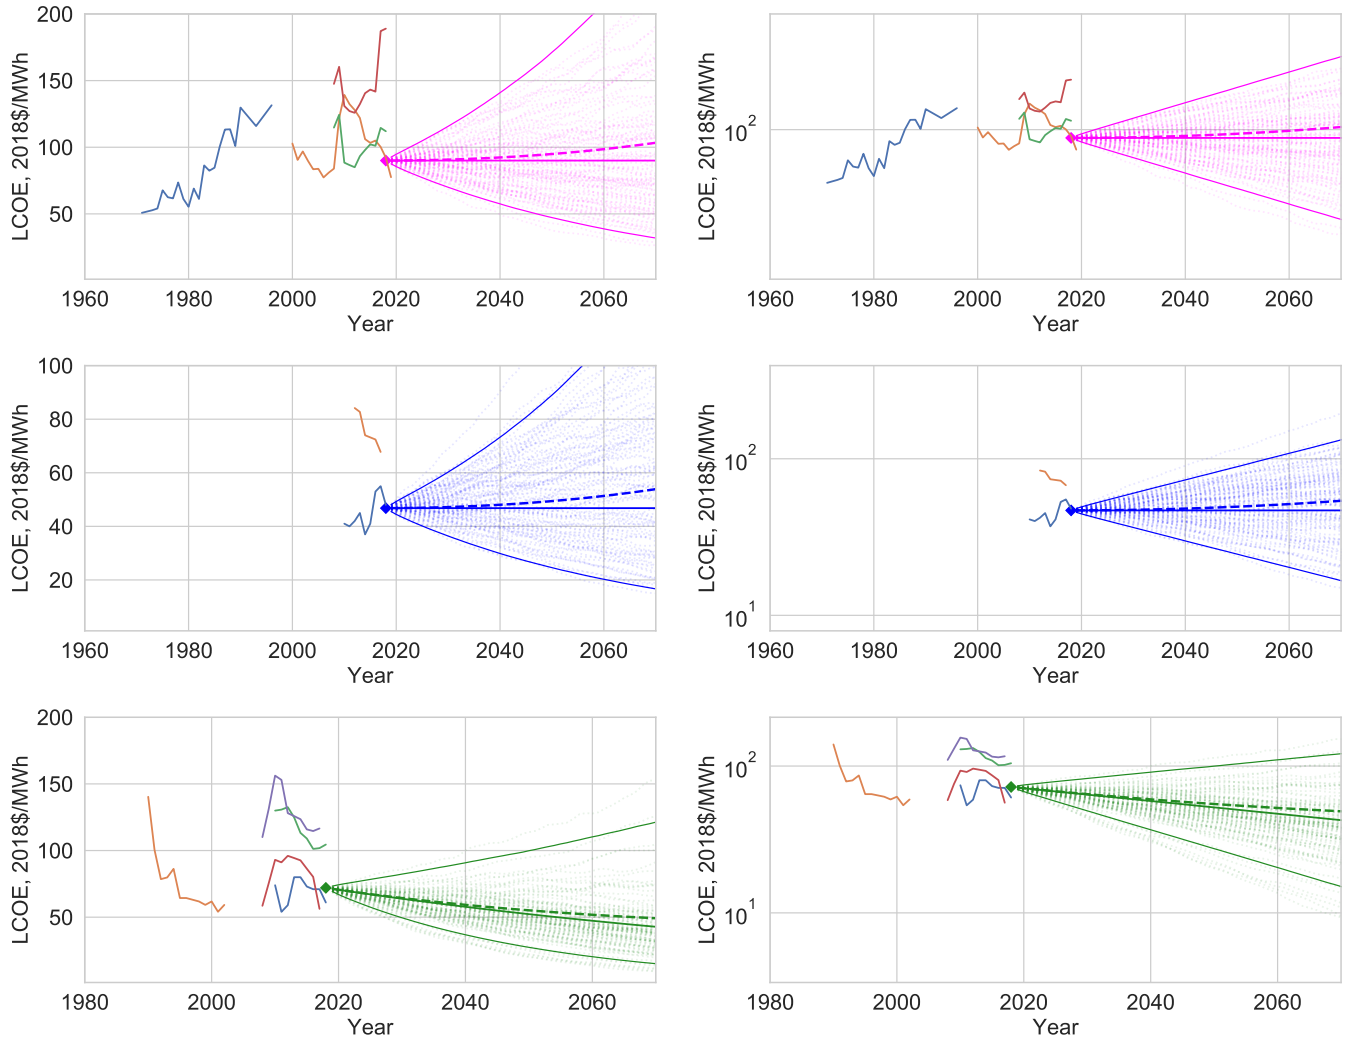

Figure S55: Moore's law cost forecast distributions for nuclear power (top), hydropower (middle) and biopower (bottom), on a non-log scale (L) and a log scale (R). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Section 6.

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 PEM PtG mean cost data used before yields parameter estimates: $\hat{\mu} = 0.143$, $\hat{\sigma}_\mu = 0.084$, $\hat{\sigma}_n = 0.251$. The data is so sparse that these values are questionable, but we use them anyway.

The FDML forecasts for batteries and electrolyzers are shown in Figure S57. Again, the forecasts beyond about 30 years begin to look unrealistic. This is especially noticeable for electrolyzers — the sparsity of the data available for calibration results in a large Moore exponent, but also very large uncertainty parameters. The effect of this is that even though the forecast median decreases, after a few decades the forecast expectation increases rapidly. This simply shows that the lack of historical data results in an uninformative forecast, and that an appropriate maximum time horizon for this model is around 2050. This inappropriate forecast feeds through directly to the system cost calculation, so great care is required. FDML parameters are summarised in Table S23.

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 S58 shows the evolution of the annual system cost distributions for the four scenarios, and Figure S59 shows the median annual spend on each technology in each scenario. Notice how the expectation of the annual system cost increases sharply after around 2050 in all three transition scenarios. This is just a consequence of the sparse data — and hence poor quality calibration — for the electrolyzer cost model. The No Transition has no electrolyzers, and

Data source	$\hat{\mu}$	$\hat{\sigma}_\mu$	$\hat{\sigma}_n$	Start year	Final year	Series length	Current LCOE
Original data							
Nemet (2006)	0.145	0.013	0.090	1957	2003	47	
Nemet PCDB	0.104	0.020	0.112	1976	2009	34	
Schilling & Esmundo (2009) (mean)	0.061	0.002	0.010	1980	2005	26	
Lazard LCOE Analysis (Low)	0.234	0.041	0.129	2009	2019	11	40 $/MWh
Lazard LCOE Analysis (High)	0.234	0.045	0.144	2009	2019	11	46 $/MWh
Bloomberg Terminal / BNEF	0.176	0.021	0.064	2009	2018	10	70 $/MWh
LBNL Levelized PPA	0.163	0.030	0.109	2006	2019	14	28.2 $/MWh
IRENA (2019)	0.184	0.026	0.074	2010	2018	9	85 $/MWh
Chained data							
Nemet PCDB and IRENA (2019)	0.119	0.017	0.109	1976	2018	43	85 $/MWh
Nemet PCDB and Bloomberg Terminal	0.124	0.017	0.109	1976	2018	43	70 $/MWh

Table S22: Moore's law parameters for solar PV (*c.f.* Table S15).

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	0.128	0.021	0.099
Va-redox flow batteries	0.076	0.015	0.086
PEM electrolyzers	0.143	0.084	0.251

Table S23: Summary of Moore's law parameters.

hence does not display the same behvaiour. At the same time, as a result of the Moore's law forecasts being generally much lower than the Wright's law forecasts for wind, solar and batteries, the median annual system cost forecasts for the Fast and Slow transition are much lower than when using Wright's law.

Figure S60 shows the expectations and medians of the NPC of each scenario relative to No Transition, for varying discount rates, at varying time horizons. Due to the extreme parameter uncertainty for electrolyzers, we present the results for time horizons from 2040-2070 only. 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, and that compared to all other scenarios, the Fast Transition scenario is expected to generate large savings at all relevant discount rates and time horizons.

References

Anandarajah, G. & McDowall, W. (2015), Multi-cluster technology learning in times: a transport sector case study with tiam-ucl, *in* 'Informing energy and climate policies using energy systems models', Springer, pp. 261–278.

Ansar, A., Flyvbjerg, B., Budzier, A. & Lunn, D. (2014), 'Should we build more large dams? the actual costs of hydropower megaproject development', *Energy Policy* **69**, 43 – 56.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421513010926*

Arrow, K. J. (1962), 'The economic implications of learning by doing', *The review of economic studies* pp. 155–173.

Barbose, G., Darghouth, N., Elmallah, S., Forrester, S., LaCommare, K., Millstein, D., Rand, J., Cotton, W., Sherwood, S. & 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.

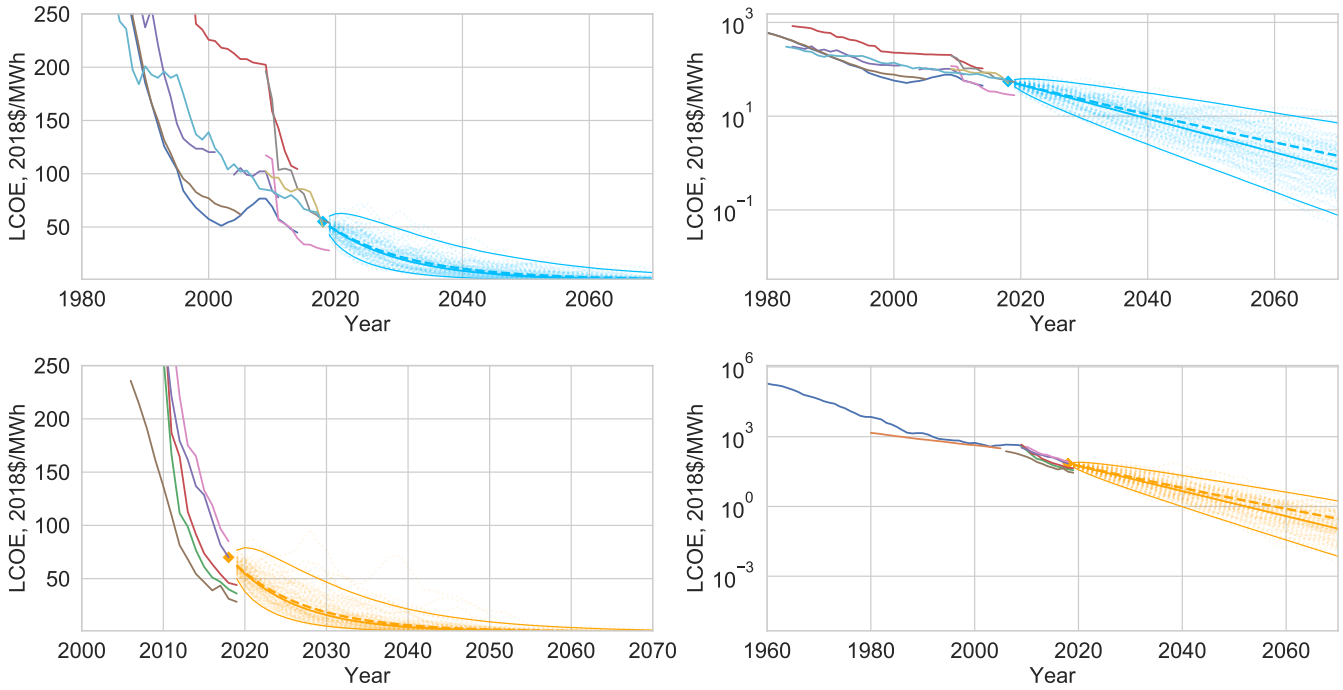

Figure S56: Moore's law cost forecast distributions for solar (top) and wind (bottom), on a non-log scale (L) and a log scale (R). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Section 6.

Berglund, C. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421504002927*

Bhattacharyya, S. C. & Timilsina, G. R. (2010), 'A review of energy system models', *International Journal of Energy Sector Management* .

Blanco, H. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S1364032117311310*

BNEF (2019), 'Electric Vehicle Outlook 2019'.

Boccard, N. (2014), 'The cost of nuclear electricity: France after fukushima', *Energy Policy* **66**, 450 – 461.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421513011440*

BP (2019), BP Statistical Review of World Energy 2019, Technical report, BP plc.

Brouwer, A. S., van den Broek, M., Zappa, W., Turkenburg, W. C. & Faaij, A. (2016), 'Least-cost options for integrating intermittent renewables in low-carbon power systems', *Applied Energy* **161**, 48 – 74.
URL: *http://www.sciencedirect.com/science/article/pii/S0306261915012167*

Brown, D., Cabe, J. & Stout, T. (2011), 'National lab uses ogj data to develop cost equations', *Oil and Gas Journal* **109**(1), 108 – 111.

Brynolf, S., Taljegard, M., Grahn, M. & Hansson, J. (2018), 'Electrofuels for the transport sector: A review of production costs', *Renewable and Sustainable Energy Reviews* **81**, 1887 – 1905.
URL: *http://www.sciencedirect.com/science/article/pii/S1364032117309358*

Butera, G., Jensen, S. H. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0360544218320693*

Buttler, A. & 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*

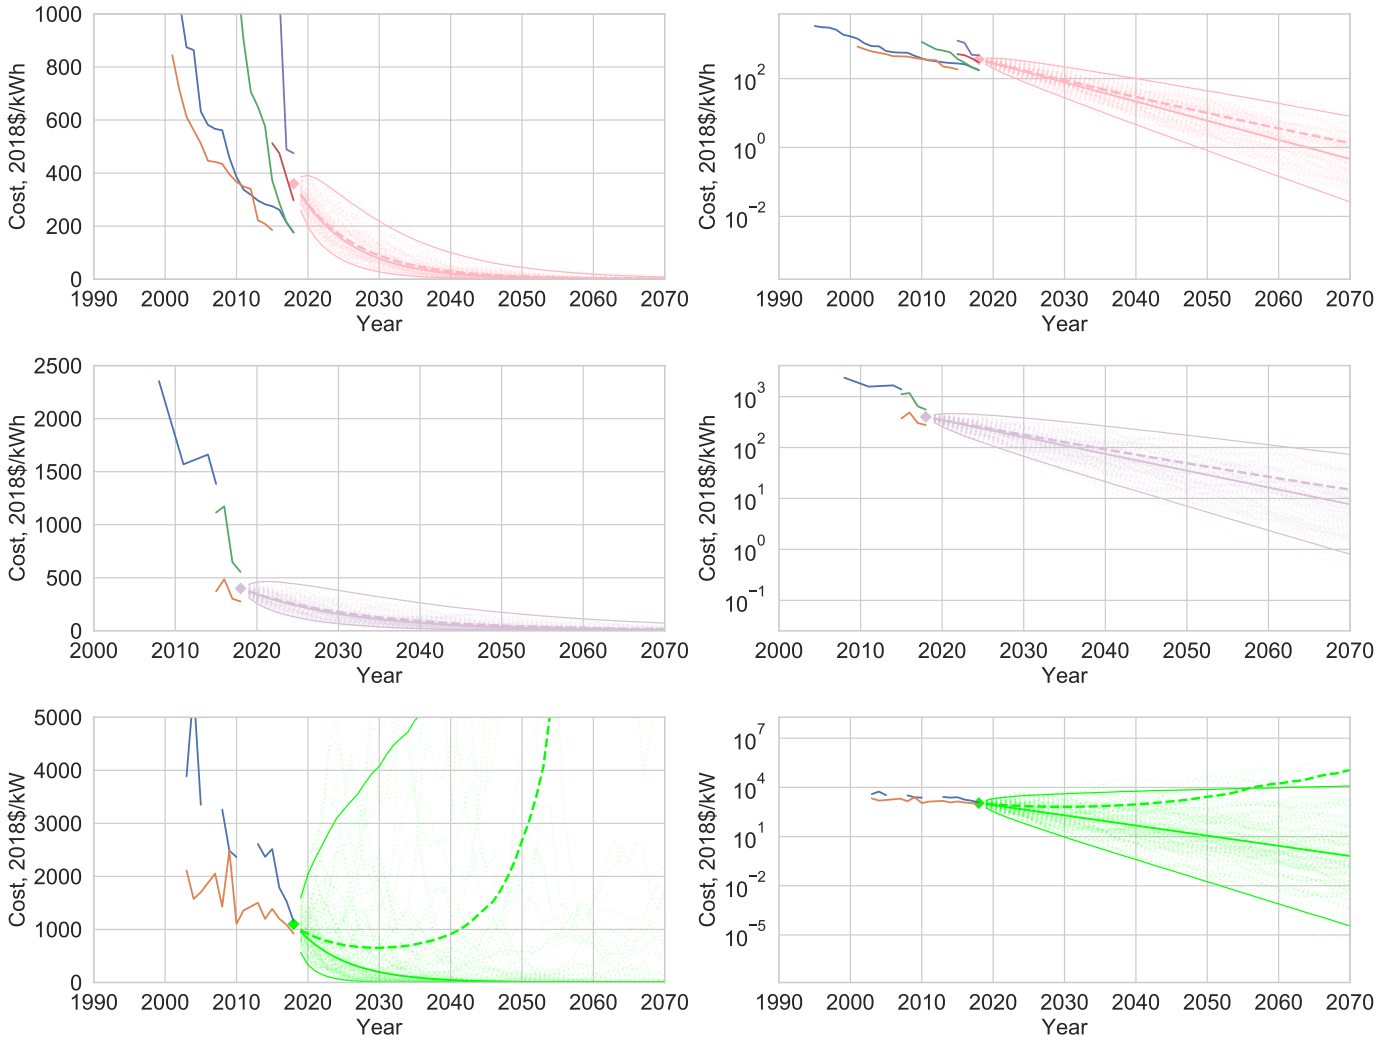

Figure S57: Moore's law cost forecast distributions for Li-ion batteries (top), Va-redux flow batteries (middle) and PEM electrolyzers (bottom), on a non-log scale (L) and a log scale (R). Cost forecast distributions show 100 sample paths, medians (solid), expectations (dashed) and 95% confidence intervals. Historical data is the same as in Section 6.

82, 2440 – 2454.
URL: *http://www.sciencedirect.com/science/article/pii/S136403211731242X*

Cannon, D., Brayshaw, D., Methven, J., Coker, P. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S096014811400651X*

Carmo, M., Fritz, D. L., Mergel, J. & Stolten, D. (2013), 'A comprehensive review on pem water electrolysis', *International Journal of Hydrogen Energy* **38**(12), 4901 – 4934.
URL: *http://www.sciencedirect.com/science/article/pii/S0360319913002607*

Citibank (2016), 'Investment Themes in 2016'.

Clark, R. (2019), 'Solbat – challenges and opportunities from an industrial (materials) perspective', Faraday Institution meeting, February 21, 2019.

Colpier, U. C. & Cornland, D. (2002), 'The economics of the combined cycle gas turbine—an experience curve analysis', *Energy Policy* **30**(4), 309 – 316.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421501000970*

Court of Audit (2012), The costs of the nuclear power sector, Thematic public report, French Court of Audit.

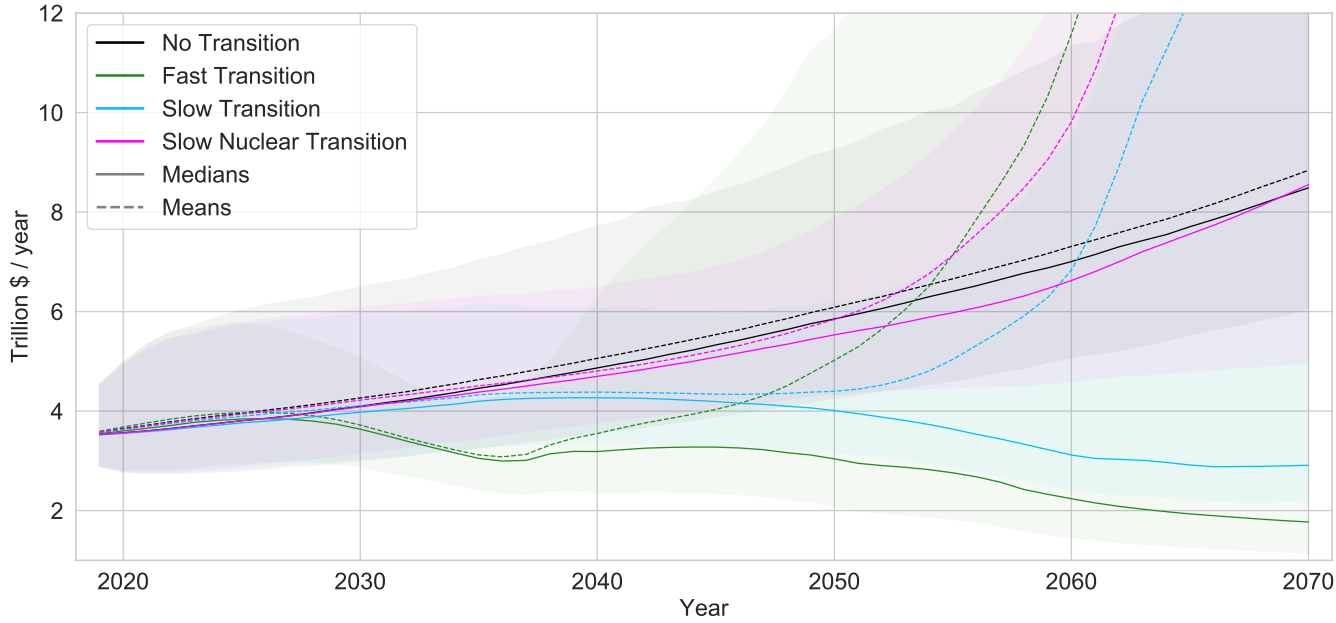

Figure S58: Annual system cost distributions, $W_t^{scenario}$, for each scenario when using the Moore's law forecasting model. Medians are shown as solid lines, expectations as dashed lines, and the shaded regions represent the 95% confidence intervals.

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. & Weber, E. U. (2018), 'Towards demand-side solutions for mitigating climate change', *Nature Climate Change* **8**(4), 260–263.
URL: *https://doi.org/10.1038/s41558-018-0121-1*

De Stercke, S. (2014), Dynamics of energy systems: A useful perspective, IIASA interim report, IIASA, Laxenburg, Austria.
URL: *http://pure.iiasa.ac.at/id/eprint/11254/*

Deason, J., Wei, M., Leventis, G., Smith, S. & Schwartz, L. C. (2018), Electrification of buildings and industry in the united states: Drivers, barriers, prospects, and policy approaches, Technical report, LBNL.

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.

Denholm, P. & Mai, T. (2019), 'Timescales of energy storage needed for reducing renewable energy curtailment', *Renewable Energy* **130**, 388 – 399.
URL: *http://www.sciencedirect.com/science/article/pii/S0960148118307316*

Dodds, P. E., Staffell, I., Hawkes, A. D., Li, F., Grünewald, P., McDowall, W. & Ekins, P. (2015), 'Hydrogen and fuel cell technologies for heating: A review', *International Journal of Hydrogen Energy* **40**(5), 2065 – 2083.
URL: *http://www.sciencedirect.com/science/article/pii/S0360319914031383*

EIA (2019*a*), Annual coal report 2018, Technical report, U.S. Energy Information Administration.

EIA (2019*b*), Annual Energy Outlook reports 1979-2019, Technical report, U.S. Energy Information Administration.

EIA (2019*c*), 'Gasoline and diesel fuel update'. [Online; accessed 31-December-2019].
URL: *https://www.eia.gov/petroleum/gasdiesel/*

Ellenbeck, S. & Lilliestam, J. (2019), 'How modelers construct energy costs: discursive elements in energy system and integrated assessment models', *Energy Research & Social Science* **47**, 69–77.

EUTurbines (2019), 'EUTurbines renewable gas commitments 2019'. [Online; accessed 22-December-2019].
URL: *https://powertheeu.eu/*

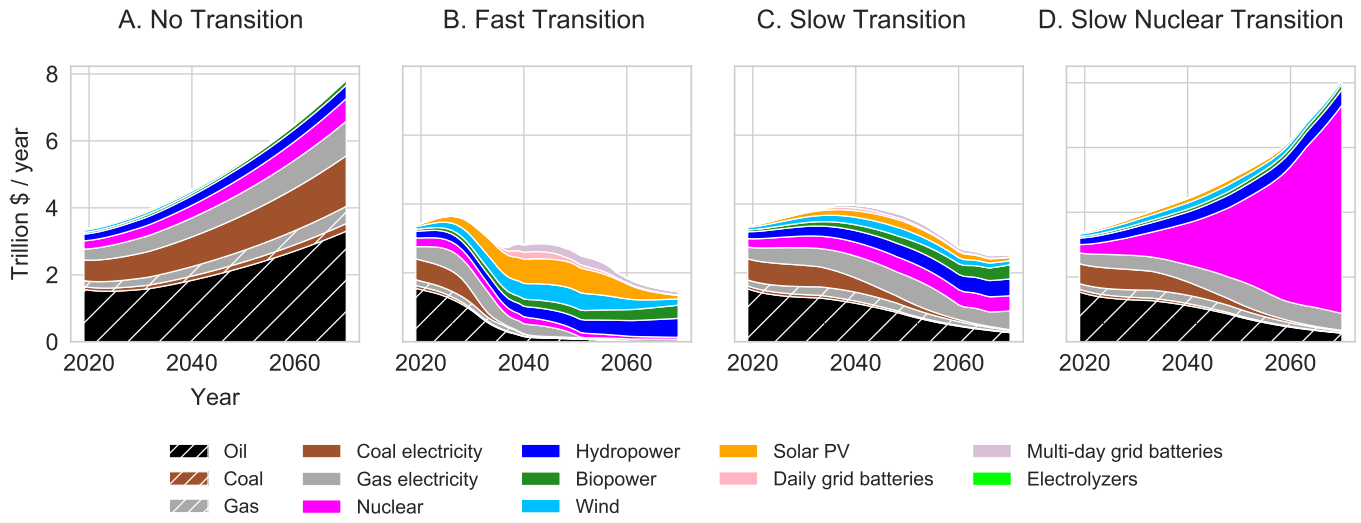

Figure S59: Median annual spend on each technology in each scenario when using the Moore's law forecasting model.

Farmer, J. D. & Lafond, F. (2016), 'How predictable is technological progress?', *Research Policy* **45**(3), 647–665.

Ferioli, F., Schoots, K. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421508006101*

Fouquet, R. (2008), *Heat, Power and Light: revolutions in energy services*, Edward Elgar Publications, Cheltenham, UK, and Northampton, MA, USA.

FRED (2019), 'Economic research at the federal reserve bank of st. louis'. [Online; accessed 31-December-2019].
URL: *https://fred.stlouisfed.org/*

Fu, R., Remo, T. W. & Margolis, R. M. (2018), '2018 u.s. utility-scale photovoltaics-plus-energy storage system costs benchmark'.

GEA (2012), *Global Energy Assessment - Toward a Sustainable Future*, Cambridge University Press, Cambridge, UK and New York, NY, USA and the International Institute for Applied Systems Analysis, Laxenburg, Austria.

Gilbert, A., Sovacool, B. K., Johnstone, P. & Stirling, A. (2017), 'Cost overruns and financial risk in the construction of nuclear power reactors: A critical appraisal', *Energy Policy* **102**, 644 – 649.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421516301690*

Glenk, G. & Reichelstein, S. (2019), 'Economics of converting renewable power to hydrogen', *Nature Energy* **4**(3), 216–222.
URL: *https://doi.org/10.1038/s41560-019-0326-1*

Grams, C. M., Beerli, R., Pfenninger, S., Staffell, I. & Wernli, H. (2017), 'Balancing europe's wind-power output through spatial deployment informed by weather regimes', *Nature Climate Change* **7**(8), 557–562.
URL: *https://doi.org/10.1038/nclimate3338*

Grubler, A. (2010), 'The costs of the french nuclear scale-up: A case of negative learning by doing', *Energy Policy* **38**(9), 5174 – 5188. Special Section on Carbon Emissions and Carbon Management in Cities with Regular Papers.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421510003526*

Grubler, A., Nakicenovic, N. & Victor, D. G. (1999), 'Dynamics of energy technologies and global change', *Energy Policy* **27**(5), 247 – 280.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421598000676*

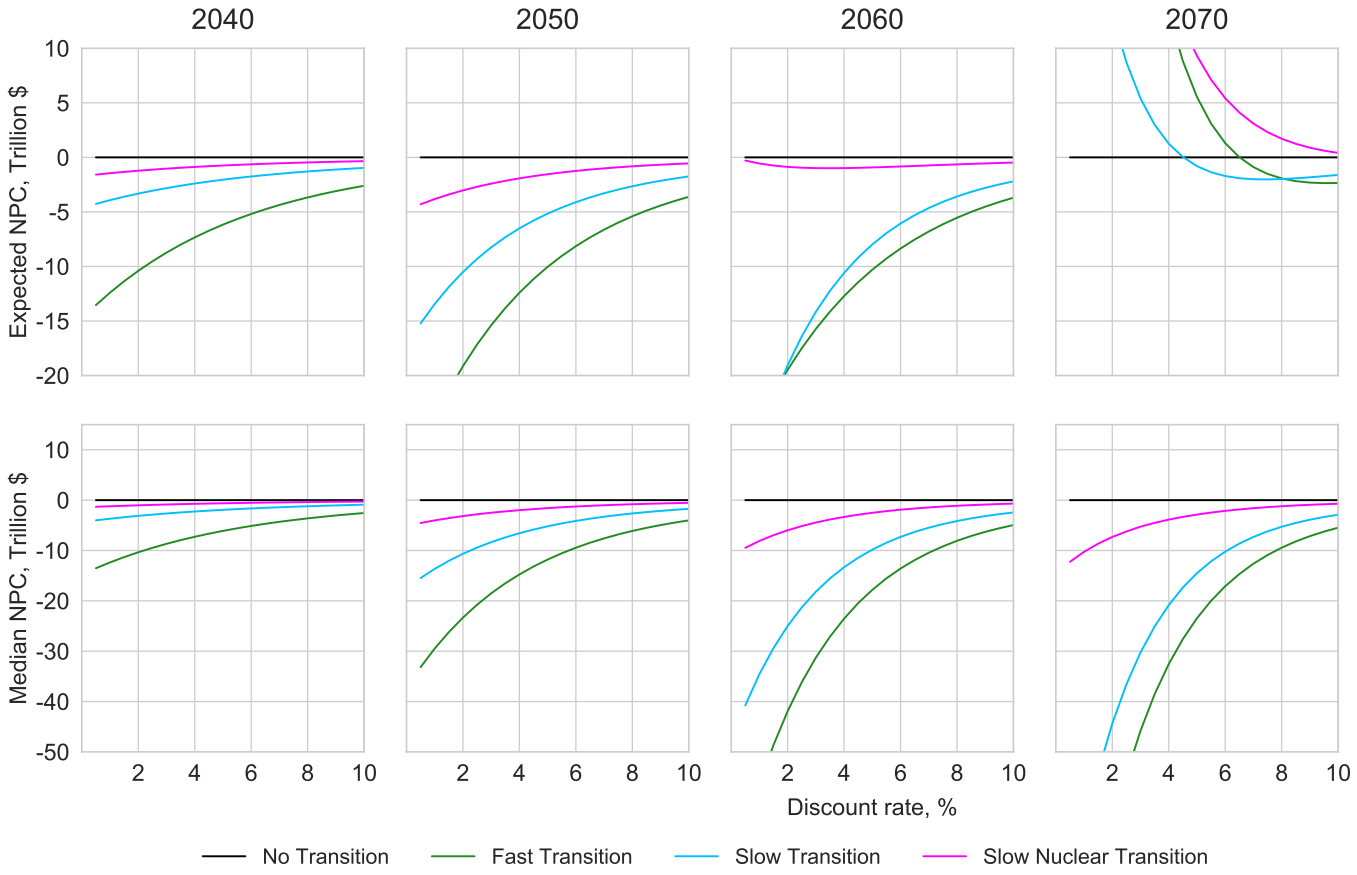

Figure S60: Expected NPC (upper row) and median NPC (lower row) relative to No Transition, against discount rate, for the four scenarios, when performing the NPC calculation using different time horizons between 2040 and 2070, and the FDML.

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. & 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.
URL: *https://doi.org/10.1038/s41560-018-0172-6*

Ha, S. & Gallagher, K. G. (2015), 'Estimating the system price of redox flow batteries for grid storage', *Journal of Power Sources* **296**, 122 – 132.
URL: *http://www.sciencedirect.com/science/article/pii/S0378775315300525*

Harding, R. (2019), 'Japan launches first liquid hydrogen carrier ship'. [Online; accessed 11-January-2020].
URL: *https://www.ft.com/content/8ae16d5e-1bd4-11ea-97df-cc63de1d73f4*

Heaps, C. (2008), 'An introduction to leap', *Stockholm Environment Institute* pp. 1–16.

Heide, D., Greiner, M., von Bremen, L. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0960148111000851*

Heuberger, C. F., Rubin, E. S., Staffell, I., Shah, N. & Mac Dowell, N. (2017), 'Power capacity expansion planning considering endogenous technology cost learning', *Applied Energy* **204**, 831–845.

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.

Icerman, L. (1974), 'Relative costs of energy transmission for hydrogen, natural gas, and electricity', *Energy Sources* **1**(4), 435–446.
URL: *https://doi.org/10.1080/00908317408945936*

IEA (2017), Energy Technology Perspectives 2017 - Catalysing Energy Technology Transformations, Technical report, International Energy Agency.

IEA (2018), Global EV Outlook 2018 - Towards cross-modal electrification, Technical report, International Energy Agency.

IEA (2019*a*), Key World Energy Statistics, Technical report, International Energy Agency.

IEA (2019*b*), The Future of Hydrogen, Technical report, International Energy Agency.

IEA (2019*c*), World Energy Outlook 2019, Technical report, International Energy Agency.

Ikäheimo, J., Kiviluoma, J., Weiss, R. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0360319918319931*

IRENA (2018*a*), Hydrogen from renewable power: Technology outlook for the energy transition, Technical report, International Renewable Energy Agency, Abu Dhabi.

IRENA (2018*b*), Renewable Power Generation Costs in 2017, Technical report, International Renewable Energy Agency, Abu Dhabi.

IRENA (2019), Renewable Power Generation Costs in 2018, Technical report, International Renewable Energy Agency, Abu Dhabi.

Jensen, S. H., Graves, C., Mogensen, M., Wendel, C., Braun, R., Hughes, G., Gao, Z. & 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.
URL: *http://dx.doi.org/10.1039/C5EE01485A*

Junginger, M., de Visser, E., Hjort-Gregersen, K., Koornneef, J., Raven, R., Faaij, A. & Turkenburg, W. (2006), 'Technological learning in bioenergy systems', *Energy Policy* **34**(18), 4024 – 4041.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421505002533*

Kassakian, J. G., Schmalensee, R., Hogan, W. W., Jacoby, H. D. & Kirtley, J. L. (2011), The future of the electric grid, Technical report, MIT.

Keramidas, K., Kitous, A., Després, J. & Schmitz, A. (2017), 'Poles-jrc model documentation', *Publications Office of the European Union* .

Kittner, N., Lill, F. & Kammen, D. M. (2017), 'Energy storage deployment and innovation for the clean energy transition', *Nature Energy* **2**(9), 17125.
URL: *https://doi.org/10.1038/nenergy.2017.125*

Kittner, N., Schmidt, O., Staffell, I. & Kammen, D. M. (2020), Chapter 8 - grid-scale energy storage, *in* M. Junginger & A. Louwen, eds, 'Technological Learning in the Transition to a Low-Carbon Energy System', Academic Press, pp. 119 – 143.
URL: *http://www.sciencedirect.com/science/article/pii/B978012818762300008X*

Kittner, N., Tsiropoulos, I., Tarvydas, D., Schmidt, O., Staffell, I. & Kammen, D. M. (2020), Chapter 9 - electric vehicles, *in* M. Junginger & A. Louwen, eds, 'Technological Learning in the Transition to a Low-Carbon Energy System', Academic Press, pp. 145 – 163.
URL: *http://www.sciencedirect.com/science/article/pii/B9780128187623000091*

Kobayashi, H., Hayakawa, A., Somarathne, K. & Okafor, E. (2019), 'Science and technology of ammonia combustion', *Proceedings of the Combustion Institute* **37**(1), 109 – 133.
URL: *http://www.sciencedirect.com/science/article/pii/S1540748918306345*

Koomey, J. & Hultman, N. E. (2007), 'A reactor-level analysis of busbar costs for us nuclear plants, 1970–2005', *Energy Policy* **35**(11), 5630 – 5642.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421507002558*

Koomey, J., Hultman, N. E. & Grubler, A. (2017), 'A reply to "historical construction costs of global nuclear power reactors"', *Energy Policy* **102**, 640 – 643.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421516301549*

Kramer, G. J. & Haigh, M. (2009), 'No quick switch to low-carbon energy', *Nature* **462**(7273), 568.

Lafond, F., Bailey, A. G., Bakker, J. D., Rebois, D., Zadourian, R., McSharry, P. & 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.

Lafond, F., Greenwald, D. & Farmer, J. D. (2020), 'Can stimulating demand drive costs down? world war ii as a natural experiment', *INET Working Paper* .

Lazard (2019*a*), Lazard's Levelized Cost of Energy ("LCOE") Analysis Versions 2-13, Technical report, Lazard.

Lazard (2019*b*), Lazard's Levelized Cost of Storage ("LCOS") Analysis Versions 1-5, Technical report, Lazard.

Li, Z., Pan, M. S., Su, L., Tsai, P.-C., Badel, A. F., Valle, J. M., Eiler, S. L., Xiang, K., Brushett, F. R. & Chiang, Y.-M. (2017), 'Air-breathing aqueous sulfur flow battery for ultralow-cost long-duration electrical storage', *Joule* **1**(2), 306 – 327.
URL: *http://www.sciencedirect.com/science/article/pii/S2542435117300326*

Loulou, R., Goldstein, G., Noble, K. et al. (2004), 'Documentation for the markal family of models', *Energy Technology Systems Analysis Programme* pp. 65–73.

Loulou, R. & Labriet, M. (2008), 'Etsap-tiam: the times integrated assessment model part i: Model structure', *Computational Management Science* **5**(1-2), 7–40.

Loulou, R., Remme, U., Kanudia, A., Lehtila, A. & Goldstein, G. (2005), 'Documentation for the times model part ii', *Energy Technology Systems Analysis Programme* .

Lovering, J. R., Yip, A. & Nordhaus, T. (2016), 'Historical construction costs of global nuclear power reactors', *Energy Policy* **91**, 371 – 382.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421516300106*

Lovins, A. B. (2017), 'Reliably integrating variable renewables: Moving grid flexibility resources from models to results', *The Electricity Journal* **30**(10), 58 – 63.
URL: *http://www.sciencedirect.com/science/article/pii/S1040619017302889*

Lovins, A. B. (2018), 'How big is the energy efficiency resource?', *Environmental Research Letters* **13**(9), 090401.
URL: *https://doi.org/10.1088%2F1748-9326%2Faad965*

Mai, T., Jadun, P., Logan, J., McMillan, C., Muratori, M., Steinberg, D., Vimmerstedt, L., Jones, R., Haley, B. & Nelson, B. (2018), Electrification futures study: Scenarios of electric technology adoption and power consumption for the united states, Nrel/tp-6a20-71500, NREL, National Renewable Energy Laboratory, Golden, CO.
URL: *https://www.nrel.gov/docs/fy18osti/71500.pdf*

Matsuo, Y. & Nei, H. (2019), 'An analysis of the historical trends in nuclear power plant construction costs: The japanese experience', *Energy Policy* **124**, 180 – 198.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421518305962*

Mazzola, J. B. & McCardle, K. F. (1997), 'The stochastic learning curve: Optimal production in the presence of learning-curve uncertainty', *Operations Research* **45**(3), 440–450.

McNerney, J., Farmer, J. D. & Trancik, J. E. (2011), 'Historical costs of coal-fired electricity and implications for the future', *Energy Policy* **39**(6), 3042 – 3054.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421511000474*

Messner, S. & Strubegger, M. (1995), 'User's guide for message iii'.

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.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421597001353*

Nemet, G. F. (2006), 'Beyond the learning curve: factors influencing cost reductions in photovoltaics', *Energy Policy* **34**(17), 3218 – 3232.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421505001795*

Nordhaus, W. D. (2014), 'The perils of the learning model for modeling endogenous technological change', *The Energy Journal* **Volume 35**(1).
URL: *http://www.iaee.org/en/publications/ejarticle.aspx?id=2546*

Nykvist, B. & Nilsson, M. (2015), 'Rapidly falling costs of battery packs for electric vehicles', *Nature Climate Change* **5**(4), 329–332.
URL: *https://doi.org/10.1038/nclimate2564*

OECD & NEA (2019), The costs of decarbonisation: System costs with high shares of nuclear and renewables, Technical report, OECD and Nuclear Energy Agency.

Ogden, J. M. (2018), Prospects for hydrogen in the future energy system, Technical report, Institute of Transportation Studies, University of California, Davis.

Pillot, C. (2018), 'Current Status and Future Trends of the Global Li-ion Battery Market'.

Pindyck, R. S. (1999), 'The long-run evolution of energy prices', *The Energy Journal* **20**(2), 1–27.
URL: *http://www.jstor.org/stable/41322828*

Pleßmann, G., Erdmann, M., Hlusiak, M. & Breyer, C. (2014), 'Global energy storage demand for a 100 8th International Renewable Energy Storage Conference and Exhibition (IRES 2013).
URL: *http://www.sciencedirect.com/science/article/pii/S1876610214001702*

Rasmussen, M. G., Andresen, G. B. & Greiner, M. (2012), 'Storage and balancing synergies in a fully or highly renewable pan-european power system', *Energy Policy* **51**, 642 – 651. Renewable Energy in China.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421512007677*

Ritchie, H. & Roser, M. (2018), 'Our world in data, energy production and changing energy sources'. [Online; accessed 22-December-2019].
URL: *https://ourworldindata.org/energy-production-and-changing-energy-sources*

Rubin, E. S., Azevedo, I. M., Jaramillo, P. & Yeh, S. (2015), 'A review of learning rates for electricity supply technologies', *Energy Policy* **86**, 198 – 218.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421515002293*

Rubin, E. S., Davison, J. E. & Herzog, H. J. (2015), 'The cost of CO_2 capture and storage', *International Journal of 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 CO_2 Capture and Storage.
URL: *http://www.sciencedirect.com/science/article/pii/S1750583615001814*

Saarinen, L., Dahlbäck, N. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0960148115003286*

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.
URL: *http://www.sciencedirect.com/science/article/pii/S1364032117312224*

Schilling, M. A. & Esmundo, M. (2009), 'Technology s-curves in renewable energy alternatives: Analysis and implications for industry and government', *Energy Policy* **37**(5), 1767 – 1781.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421509000111*

Schlömer, S., Bruckner, T., Fulton, L., Hertwich, E., McKinnon, A., Perczyk, D., Roy, J., Schaeffer, R., Sims, R., Smith, P. & Wiser, R. (2014), Annex iii: Technology-specific cost and performance parameters, *in* O. Edenhofer, R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel & J. Minx, eds, '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, Cambridge, United Kingdom and New York, NY, USA.

Schmidt, O., Hawkes, A., Gambhir, A. & Staffell, I. (2017), 'The future cost of electrical energy storage based on experience rates', *Nature Energy* **2**(8), 17110.
URL: *https://doi.org/10.1038/nenergy.2017.110*

Schoots, K., Ferioli, F., Kramer, G. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0360319908002954*

Seebregts, A. J. (2010), Gas-fired power, Technology brief e02, IEA ETSAP.

Shafiee, S. & Topal, E. (2010), 'A long-term view of worldwide fossil fuel prices', *Applied Energy* **87**(3), 988 – 1000.
URL: *http://www.sciencedirect.com/science/article/pii/S0306261909003912*

Shaner, M. R., Davis, S. J., Lewis, N. S. & Caldeira, K. (2018), 'Geophysical constraints on the reliability of solar and wind power in the united states', *Energy Environ. Sci.* **11**, 914–925.
URL: *http://dx.doi.org/10.1039/C7EE03029K*

Smil, V. (2016), *Energy Transitions, Global and National Perspectives, 2nd Edition*, ABC-CLIO, Santa Barbara, CA.

Solomon, A., Kammen, D. M. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S0306261914007867*

Sovacool, B. K., Gilbert, A. & Nugent, D. (2014), 'Risk, innovation, electricity infrastructure and construction cost overruns: Testing six hypotheses', *Energy* **74**, 906 – 917.
URL: *http://www.sciencedirect.com/science/article/pii/S0360544214008925*

Speirs, J., Balcombe, P., Johnson, E., Martin, J., Brandon, N. & Hawkes, A. (2018), 'A greener gas grid: What are the options', *Energy Policy* **118**, 291 – 297.
URL: *http://www.sciencedirect.com/science/article/pii/S0301421518302027*

Staffell, I., Brett, D., Brandon, N. & Hawkes, A. (2012), 'A review of domestic heat pumps', *Energy Environ. Sci.* **5**, 9291–9306.
URL: *http://dx.doi.org/10.1039/C2EE22653G*

Staffell, I., Scamman, D., Abad, A. V., Balcombe, P., Dodds, P. E., Ekins, P., Shah, N. & Ward, K. R. (2019), 'The role of hydrogen and fuel cells in the global energy system', *Energy & Environmental Science* **12**(2), 463–491.

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.

Trancik, J. E., Jean, J., Kavlak, G., Klemun, M. M., Edwards, M. R., McNerney, J., Miotti, M., Brown, P. R., Mueller, J. M. & Needell, Z. A. (2015), Technology improvement and emissions reductions as mutually reinforcing efforts: Observations from the global development of solar and wind energy, Technical report, MIT.

Trutnevyte, E. (2016), 'Does cost optimization approximate the real-world energy transition?', *Energy* **106**, 182–193.

Vaillancourt, K. (2014), Coal mining and logistics, Technology brief p07/08, IEA ETSAP.

van der Zwaan, B., Schoots, K., Rivera-Tinoco, R. & 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.
URL: *http://www.sciencedirect.com/science/article/pii/S030626191100314X*

Way, R., Lafond, F., Lillo, F., Panchenko, V. & 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* .

WEC (2018), Bringing north sea energy ashore efficiently, Technical report, Word Energy Council.

Wene, C.-O. (2016), 'Future energy system development depends on past learning opportunities', *WIREs Energy and Environment* **5**(1), 16–32.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/wene.172*

Wilson, C., Grubler, A., Bento, N., Healey, S., De Stercke, S. & Zimm, C. (2020), 'Granular technologies to accelerate decarbonization', *Science* **368**(6486), 36–39.
URL: *https://science.sciencemag.org/content/368/6486/36*

Wilson, C., Grubler, A., Gallagher, K. S. & Nemet, G. F. (2012), 'Marginalization of end-use technologies in energy innovation for climate protection', *Nature Climate Change* **2**(11), 780–788.
URL: *https://doi.org/10.1038/nclimate1576*

Wiser, R., Jenni, K., Seel, J., Baker, E., Hand, M., Lantz, E. & Smith, A. (2016), 'Expert elicitation survey on future wind energy costs', *Nature Energy* **1**(10), 16135.
URL: *https://doi.org/10.1038/nenergy.2016.135*

Zeyringer, M., Price, J., Fais, B., Li, P.-H. & 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.
URL: *https://doi.org/10.1038/s41560-018-0128-x*

Ziegler, M. S., Mueller, J. M., Pereira, G. D., Song, J., Ferrara, M., Chiang, Y.-M. & Trancik, J. E. (2019), 'Storage requirements and costs of shaping renewable energy toward grid decarbonization', *Joule* **3**(9), 2134 – 2153.
URL: *http://www.sciencedirect.com/science/article/pii/S2542435119303009*