

The impact of the Net Zero transition on aggregate productivity

Emilien Ravigné^{1,2,3,4} and François Lafond^{1,2}

¹*Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford, UK*

²*Smith School of Enterprise and the Environment, University of Oxford, Oxford, UK*

³*CMCC Foundation - Euro-Mediterranean Center on Climate Change, Italy*

⁴*RFF-CMCC European Institute on Economics and the Environment, Italy*

Abstract

The net-zero transition requires investment but is expected to induce cost declines in various parts of the economy, notably in renewable energy. The overall effect on aggregate productivity growth, a key driver of GDP and income, is therefore an empirical question. Here, we develop a simple framework that maps sectoral price and quantity paths, readily available from techno-economic transition scenarios, into aggregate total factor productivity (TFP) growth. Our approach combines a reduced-form industry-level mapping and Domar aggregation. We validate both steps on historical data, to show that our long-horizon forecasts are close to unbiased and to use backtest forecast errors as data-driven prediction intervals. We apply the framework UK economy-wide scenarios and find small aggregate effects by 2050, either slightly positive or modestly negative (-1.3% to +0.1%). By contrast, global energy transition scenarios with rapid renewable deployment generate substantial aggregate productivity gains, on the order of 6% by 2050.

Keywords: Aggregate productivity; Net-zero transition

Significance Statement. The net-zero transition will reshape many sectors, but it does not mechanically translate into lower economy-wide productivity. Sector-level shocks propagate and combine through the existing production network. This network structure is slow to evolve, which makes the aggregate consequences of transition scenarios more constrained and more predictable than commonly assumed. Using historical data, we show that a simple mapping from scenario price and quantity paths to aggregate TFP performs well out of sample. Results show that aggregate outcomes are often less pessimistic than sector-level narratives imply, because sectoral cost increases translate into aggregate productivity losses only when they affect sufficiently important sectors, and declining renewable costs can generate positive productivity effects that offset losses elsewhere.

Introduction

The net zero transition is a long-term, economy-wide transformation [1], suggesting that it has the scale to become either a drag on economic performance or a source of green growth. While there exist detailed techno-economic scenarios showing how to achieve the transition, there is relatively little work on assessing their macroeconomic consequences. Net-zero scenarios usually contained detailed projections for technology deployment, and often report direct sectoral costs and benefits [2–4], typically for renewables, clean transport, clean heat, and clean industry options. However, there is no clear, accepted method to translate these scenario into macroeconomic numbers. Existing studies find different effects on TFP, labour productivity, and wages, either from sectoral or firm-level evidence that does not aggregate to an economy-wide net-zero productivity estimate [5, 6], or from structural models centred on specific mechanisms such as carbon tax, or investments in the electricity sector and mobility [7–10]. Therefore, we lack transparent evidence of how bottom-up technology assumptions and scenarios translate into aggregate productivity and economic growth. Here we propose a straightforward method to find the aggregate productivity projections implied by published bottom up scenarios.

We show that in an accounting framework where the production network structure is explicitly taken into account, and with minimal additional assumptions, the aggregate productivity effects of net-zero transitions can be approximated from only three objects: sectoral price changes, sectoral quantity changes, and sectors’ size, which are all typically provided by bottom-up scenarios. This result is possible using a transparent, reduced-form map between scenario-based price and production pathways and industry-level total factor productivity (TFP) and by a standard aggregation result translating industry TFP growth, with weights reflecting both their size and their position in the production network, into aggregate TFP growth. This provides a bridge from bottom-up transition scenarios to aggregate productivity without requiring a full future input-output table or a fully specified general-equilibrium model.

To validate our framework, we use backtesting on real historical data. Our method successfully recovers past productivity changes from observed price and output paths, yielding a mean forecast error close to zero. We use these historical prediction errors to calibrate empirical prediction intervals, providing a rigorous statistical foundation for our net-zero projections.

We apply the framework to two sets of scenarios. First, we evaluate the six UK Climate Change Committee (CCC) scenarios from the Sixth and Seventh Carbon Budget [4, 11], which are economy-wide and cover all industries that need to transform for a net-zero transition. We find small effects on productivity, ranging from +0.1% to -1.3% cumulative by 2050. This is not because the transition leaves sectoral productivity unchanged, but because positive and negative sectoral effects partly offset one another and because many of the most affected sectors carry limited economic weight. Second, we focus on a detailed scenario for the global energy sector [12]. Here, improvements in energy-sector productivity translate into sizeable aggregate benefits, on the order of 6% from 2022 to 2050. Applied to the UK alone, the benefits are more muted, largely due to more favourable initial conditions and less promising renewable resources.

A simple bridge from transition scenarios to aggregate productivity

Net-zero scenarios are typically expressed as future prices and quantities for specific products, technologies, or activities—for example electricity generation, electric vehicles, heat pumps, steel, or cement—rather than in terms of projections of aggregate productivity. Translating such scenarios into aggregate total factor productivity (TFP) is therefore not straightforward, for two reasons. First, bottom-up scenarios are typically constructed from emission sectors, such as energy, heat or transport, which do not align well with the standard industrial classifications used in available macroeconomic data. Second, because productivity growth is defined as the growth of output minus the growth of inputs, evaluating overall productivity growth requires tracking down output and inputs for each ladder of the value chain. For instance, if scenarios report that solar panels will become cheaper, we cannot automatically translate this into productivity growth of the solar panel manufacturing industry, as the lower costs of solar panels may be due to productivity growth in one or more of its inputs. Tracking down each industry’s inputs to assign productivity at the right level before re-aggregating would be a daunting task. Our framework clarifies these conceptual issues, and provides a simple bridge that translates transition scenarios directly into a key macroeconomic metric (Fig. 1A).

The bridge proceeds in two steps. We first map scenario-implied changes in prices and quantities into industry-level TFP growth. We then aggregate those industry contributions using Domar aggregation based on national accounting [13, 14].

Framework

Our goal is to find a function that maps τ -years ahead prices $p_{i,t,\tau}$ and quantities $X_{i,t,\tau}$ of various industries, as given by scenarios, into aggregate productivity growth between t and $t + \tau$, $\hat{A}_{t,\tau}$. We consider a closed economy, where national accounting identities hold, homogeneous goods within each sector, and short time periods.

The first step starts from a simple accounting identity: sales (i.e. nominal output) is the sum of price growth and quantity growth. We then assume that nominal input expenditure grows as a constant fraction ζ of sales growth. The SI provides a discussion of this assumption; in essence, a fixed $0 < \zeta < 1$ means that productivity growth translates into a given, specific mix of faster output growth and lower output prices. This is a testable assumption, and we will report an extensive empirical evaluation below. This assumption is extremely powerful for us: in Materials and Methods, we show that industry-level TFP growth can then be expressed as

$$\hat{A}_{i,t,\tau} = \zeta \hat{X}_{i,t,\tau} + (\zeta - 1) \hat{p}_{i,t,\tau}, \quad (1)$$

where $\hat{A}_{i,t,\tau}$ denotes TFP growth in industry i between t and $t + \tau$, $\hat{X}_{i,t,\tau}$ denotes output growth, $\hat{p}_{i,t,\tau}$ denotes output price growth. Note that Eq. 1 would be an identity if we allowed ζ to be indexed by i, t and τ - the real assumption is that ζ is a universal parameter that holds for all industries and time periods.

Once we have industry-level TFP, aggregate productivity follows from a standard result known as Domar aggregation

$$\hat{A}_{t,\tau} = \sum_i \lambda_{i,t} \hat{A}_{i,t,\tau}, \quad \lambda_{i,t} \equiv \frac{p_{i,t} X_{i,t}}{p_t Y_t}, \quad (2)$$

where $\lambda_{i,t}$ is the Domar weight of industry i , measured as industry sales $p_{i,t} X_{i,t}$ relative to GDP $p_t Y_t$. Note that Domar weights sum up to more than one, since gross output (total sales) is greater than GDP (value added). The fundamental reason for this is simple: industry-level TFP shocks are computed on gross output, while aggregate TFP is computed on value added (GDP). This second step makes it clear that TFP shocks to larger sectors have larger macroeconomic impact.

Our key result, combining Eq. 1 and 2 reads

$$\text{TFP gr.} \approx \underbrace{\sum_i \frac{\text{Sales}_{i,t}}{\text{GDP}_t}}_{\text{Domar weights}} \times \underbrace{\left(\zeta \text{Quantity gr.} + (\zeta - 1) \text{Price gr.} \right)}_{\text{Predicted industry TFP growth}}, \quad (3)$$

This bridge is the main methodological contribution of the paper. It provides a tractable way to translate heterogeneous transition scenarios into aggregate productivity outcomes while relaxing two requirements that would otherwise make the exercise infeasible: exact alignment with standard industrial classifications and complete knowledge of future input-output structures. A key assumption is that all industries on which the net-zero transition has a TFP impact are covered by the scenarios. Next, we assess the accuracy of this formula in practice.

Historical validation

We evaluate the framework by backtesting each step on historical data. The logic is straightforward. If future prices and quantities had been known at some point in the past, how good would the forecasts of industry-level TFP have been using Eq. 1? And if future industry TFP growth had been known, how good would have been the forecasts of aggregate TFP using Domar aggregation (Eq. 2). We answer these questions using three benchmark industry datasets—KLEMS, STAN, and NBER CES. This exercise both disciplines the calibration of ζ and provides the empirical error distribution used later to construct prediction intervals, in the spirit of [15–17].

At the industry level, the reduced-form mapping tracks historical manufacturing TFP reasonably well. We have determined our baseline calibration, $\zeta = 0.13$, by minimizing the absolute forecast error in the benchmark dataset. The forecast errors remain substantial for individual industries, but they are centred close to zero over long horizons (Fig. 1B).

We next assess Domar aggregation, which is not trivial over long horizons, in fact providing an empirical result of independent interest. We find that long-horizon forecast errors are again centred close to zero (Fig. 1C), although precision varies across countries. Overall, the aggregation step does not introduce large systematic distortions at the horizons relevant for the net-zero transition. As in the industry-level backtest, forecast uncertainty widens with the horizon, as



Figure 1: **A two-step bridge from sectoral transition scenarios to aggregate total-factor productivity (TFP), and its historical backtest.** (A) Framework. Scenario-implied sectoral prices $p_{i,t,\tau}$ and quantities $X_{i,t,\tau}$ are mapped to industry-level TFP growth via the sufficient-statistics relation $\hat{A}_{i,t,\tau} = \zeta \hat{X}_{i,t,\tau} + (\zeta - 1) \hat{p}_{i,t,\tau}$ (Eq. 1), then aggregated with Domar weights $\lambda_{i,t} \equiv p_{i,t} X_{i,t} / (p_t Y_t)$ (Eq. 2) to yield economy-wide TFP growth $\hat{A}_{t,\tau}$. The single free parameter ζ (share of nominal input expenditure that scales with sales) is calibrated on historical data, which also supplies the empirical error distribution used for prediction intervals. (B) Industry-level backtest: distribution of predicted-minus-observed cumulative TFP growth as a function of forecast horizon τ (years), across all manufacturing industries in EU-KLEMS, STAN, and NBER-CES, 1970–2019. (C) Aggregate-level backtest: same quantity computed after Domar aggregation, pooled across all countries. In (B, C) the solid line is the mean error, the dark and light ribbons are the 50% and 90% empirical prediction intervals, and the dashed horizontal line marks unbiased prediction. Calibration: $\zeta = 0.13$, chosen to minimise mean absolute error in (B).

expected.

The purpose of this exercise is not just to test our assumptions, but to obtain realistic uncertainty bounds usable in our own out-of-sample predictions. We use the historical distribution of industry-level forecast errors and aggregation errors to construct prediction intervals around our aggregate TFP forecasts. These intervals are data-driven, but they should be interpreted narrowly: they quantify the uncertainty introduced by our translation framework, which conditions on the underlying scenario paths for prices and quantities. Our framework does not address the accuracy or uncertainty in the scenarios themselves, it simply translates them to a key macroeconomic

quantity.

Having established that both the industry mapping and the aggregation step are approximately unbiased in historical data, we next apply the framework to UK net-zero scenarios.

Why large sectoral shocks can have negligible aggregate effects

Before turning to economy-wide scenarios, it is useful to illustrate a key implication of the framework: even dramatic cost increases in a single sector need not move aggregate productivity if that sector carries a small Domar weight. This also demonstrates the modularity of the approach — any cost assumption can be swapped in and its aggregate consequence immediately evaluated.

Cement provides a natural test case. It is a hard-to-abate, emissions-intensive sector whose decarbonisation costs feature prominently in transition debates. For each scenario, we take the implied increase in cement production costs, $\hat{p}$ to reach net zero in the cement industry in the UK by 2050, map it into a sectoral productivity effect for cement, $\hat{A}_{\text{cement}}$, using the sufficient-statistics relationship developed above, and then aggregate it using the sector’s Domar weight, 0.002 in the UK, so that $\hat{A} \approx \lambda_{\text{cement}} \hat{A}_{\text{cement}}$. Table 1 reports this calculation for three benchmark sources—IEA [18], MPP [19], and McKinsey [20]—with the scenario-implied cost increase shown explicitly in the second column. To express magnitudes in annual terms, we divide the cumulative 2022–2050 effect by 28, which provides a transparent approximation at these very small orders of magnitude.

We then stress-test the result by replacing each assumption with a deliberately extreme alternative: a much larger cost shock ($\hat{p} = 240\%$), a doubling of cement output ($\hat{X} = 100\%$), and a threefold increase in the Domar weight ($\lambda_{\text{cement}} = 0.006$). Even under these extreme assumptions, the implied aggregate effect remains very small. If the baseline TFP growth is 1% per year, then with the transition of the cement sector it would be between 0.986% and 0.999% (Table 1, last column). Although cement is an important intermediate input and a 240% price increase would raise costs for downstream users, the sector’s small economy-wide weight limits its aggregate impact.

This logic extends directly to the economy-wide analysis that follows. The CCC scenarios combine many such sectoral effects, and the aggregate result depends on their Domar-weighted sum rather than on any single sector’s cost increase.

The UK economy under Climate Change Committee (CCC) scenarios

The UK Climate Change Committee (CCC) is the government’s independent statutory adviser on climate policy and publishes detailed sectoral pathways to net zero through successive carbon budgets. Its Sixth Carbon Budget covers 2033–2037, while its later Seventh Carbon Budget covers 2038–2042.

Table 1: **Even extreme assumptions for cement imply small aggregate TFP effects.** Effects are defined in percentage points.

Scenario / assumption	Cost increase $\hat{p}$	Sectoral TFP growth	Cumulative aggregate TFP effect by 2050	Average annual effect
IEA benchmark	60%–110%	–48% to –34%	–0.13% to –0.08%	–0.004 to –0.003 %/yr
MPP benchmark	40%–120%	–50% to –25%	–0.14% to –0.06%	–0.005 to –0.002 %/yr
McKinsey benchmark	45%	–28%	–0.06%	–0.002 %/yr
Higher cost increase	240%	–66%	–0.21%	–0.007 %/yr
Output doubles ($\hat{X} = 100\%$)	40%–120%	–45% to –18%	–0.12% to –0.04%	–0.004 to –0.001 %/yr
Higher Domar weight ($\lambda_{\text{cement}} = 0.006$)	40%–120%	–48% to –34%	–0.41% to –0.18%	–0.014 tp –0.006%/yr

Using the Climate Change Committee data in our framework

The CCC scenarios [4, 11] provide data on “additional resource costs”, for each of their 14 sectors and 6 scenarios. The CCC cost data are additional to a baseline without further climate policy. Thus, the difference to that baseline captures the extra resource cost of achieving net zero, and as such our analysis is a comparison against a reference path, not a statement that productivity would be constant in the absence of the transition. Costs are free of double counting, computed at “constant service levels” (so we assume $\hat{X}_i = 0$), and annualised; see SI for details.

Let C denote the annualised extra cost of the transition at time t . The *additional* annualised extra cost between t and $t + \tau$ is $\Delta C_{i,t,\tau} = C_{i,t+\tau} - C_{i,t}$. From this, assuming price growth equals unit cost growth, we can derive the price growth between t and $t + \tau$ as $\hat{p}_{i,t,\tau} = \frac{\Delta C_{i,t,\tau}/X_i}{p_{i,t}}$, where X_i is the number of units, and p_i the price (noting that this defines $\hat{p}$ as a proper growth rate, rather than as the log growth rates we have used in our derivations, so this is only an approximation). Using this price growth (and zero quantity growth) in our framework (Eq. 3), we get $\hat{A}_{i,t,\tau} = (\zeta - 1) \frac{\Delta C_{i,t,\tau}}{p_{i,t} X_i}$. Domar aggregation then leads to an elegant solution

$$\hat{A}_{t,\tau} = (\zeta - 1) \frac{\sum_i \Delta C_{i,t,\tau}}{\text{GDP}_t}, \quad (4)$$

which implies that aggregate productivity growth between t and $t + \tau$ is equal to the total difference in (annualised, processed) costs between t and $t + \tau$, relative to GDP of the base period, and multiplied by $(\zeta - 1)$.

Results

Fig. 2 shows the time series profile of cumulative aggregate productivity growth in each of the 7th carbon budget (CB7) Balanced Pathway scenario and the 5 CCC 6th Carbon Budget (CB6) scenarios. This comparison is not fully like-for-like (see SI): CB6 annualised costs embed measure-level cost-of-capital assumptions, whereas CB7 excludes private financing costs and uses the Green Book 3.5% rate [21] only as a social discount rate for annualising capital expenditure.

The CB7 scenario brings a small positive effect of 0.1 percentage point (pp), while the CB6

scenarios show overall a small negative effect: in the worst scenario, productivity would be 1.3pp lower in 2050 compared to now. These numbers represent a very small impact on annual productivity growth rates, 0.003–0.043pp/year. The implied effects of all scenarios are not statistically distinguishable from zero, as per the uncertainty bands, and likely not even within measurement error (estimated around 0.2pp/year in the US in the 2005-2015 decade [22]).

To understand what drives these results, Fig. 2 shows the contributions of each sector to the 2022-2050 productivity growth. These results illustrate the paper’s central point. The net-zero transition is not costless at the sector level — residential buildings alone account for a cumulative drag of roughly 0.7 percentage points in some scenarios. But the aggregate effect is small because positive and negative sectoral contributions partly offset one another. Transport and electricity, for instance, contribute positively as cheaper clean technologies replace costlier incumbents, while residential buildings and CO₂ removals push in the opposite direction. The net effect ranges from slightly positive to modestly negative, and in all scenarios the prediction intervals overlap with zero, meaning the results are small.

Global and national scenarios for the energy sector

The energy production sector is likely to be the single most important industry in the net-zero transition, and some studies have found large system cost declines over the transition due to the predicted continued cost decline of renewables [12], leading to moderate macroeconomic gains, for instance aggregate real wages rising 2–3% [10] or labour productivity growing by 0.1% after 5 years [8].

Mapping energy portfolios in our framework

Because different energy sources have very different conversion losses, we express all results in terms of *useful energy*, which focuses on energy as a service. It makes it possible to compare energy sources with different conversion losses between primary energy and end use (see SI for a full discussion). This is particularly relevant for the energy transition, as renewable energy is much more efficient than fossil fuels, which waste a lot of energy as heat (e.g., in engines or power generation). While most scenarios see final energy demand drop substantially due to electrification, the amount of useful energy typically does not change dramatically.

This lets us aggregate across sources: we construct industry-level price and quantity indices from the prices and delivered quantities of each energy source in a given portfolio. We take the scenarios as given, without taking a stance on their plausibility. Unlike the simplifying case in the previous section, we no longer set quantity growth to zero; changes in the energy mix shift both the price and the quantity of useful energy delivered. We explicitly account for the cost of energy storage to address the intermittency of renewables: short- and medium-term storage is managed with batteries, while long-term storage relies on hydrogen produced from solar and wind power.

For global scenarios we take costs and energy mix from Way et al. [12]. For the UK we also use Way et al. costs, but compare two portfolios: one based on the CCC CB6 Balanced Net Zero




Figure 2: UK net-zero scenarios imply an aggregate TFP effect indistinguishable from zero, despite large sectoral contributions of opposite sign. (A) Cumulative aggregate TFP growth relative to 2022, in percentage points, for the UK Climate Change Committee (CCC) Seventh Carbon Budget Balanced Pathway (CB7; 2025) and the five Sixth Carbon Budget pathways (CB6; 2020: Balanced, Headwinds, Widespread Engagement, Widespread Innovation, Tailwinds). Shaded bands are 90% empirical prediction intervals derived from the historical backtest in Fig. 1B–C; the dashed horizontal line at zero marks the no-transition counterfactual used by CCC. By 2050, CB7 implies +0.1 pp and the CB6 scenarios range from –1.3 to –0.1 pp—equivalent to 0.003–0.043 pp yr⁻¹. (B) Decomposition of the 2022–2050 cumulative effect into the 14 CCC sectors, summed using 2022 UK Domar weights; bars above zero are productivity-raising contributions (e.g., transport, electricity), bars below zero are productivity-reducing (e.g., residential buildings, CO₂ removals). The small aggregate total arises from cancellation across sectors, not from small sectoral shocks. (C) Cumulative sectoral contributions to aggregate TFP growth relative to 2022 for the 6 scenarios of CB6 and CB7.

pathway¹, and one more renewables-heavy alternative, constructed by taking the Way et al. global mix and adjusting it to UK conditions: matching the observed 2022 UK energy mix as a starting

¹The CB7 dataset reports electricity cost inputs at the technology level but does not provide the corresponding generation mix, preventing construction of a consistent system-wide unit cost of electricity.

Table 2: Cumulative productivity growth in the energy sector, 2022–2050 (percent change, with additive log growth in parentheses).

Scenario	Price $\sum_t (\zeta - 1) \hat{p}_t$	Output $\sum_t \zeta \hat{X}_t$	Total
<i>World</i>			
Fast transition	121% (0.79)	10% (0.09)	143% (0.89)
No transition	25% (0.22)	8% (0.08)	35% (0.30)
Slow transition	56% (0.44)	9% (0.08)	69% (0.53)
<i>UK</i>			
Balanced Net Zero Pathway	15% (0.14)	-2% (-0.02)	13% (0.12)
Fast transition	64% (0.49)	-2% (-0.02)	61% (0.48)

point and replacing excess solar with wind to reflect the UK’s comparative advantage in wind generation. To compute the aggregate impact we need the Domar weight of the energy sector, which we obtain by multiplying the 2022 price of useful energy by the quantity delivered (in MWh). The resulting weight is 0.024 for the UK but nearly three times higher for the world (0.067), reflecting the larger share of energy expenditure in global GDP.

Results

Fig. 3 shows the time path of the results, and Table 2 summarises the cumulative impact and its sources. For the world (top panel), we find large energy productivity growth, 143%, driven by large price declines, moderate quantity growth, and a large Domar weight of 0.067, leading to an aggregate impact of around 5.8pp from 2022 to 2050, that is about 0.2 additional percentage points per year in TFP growth globally.

For the UK, the scenario based on the “fast transition” energy mix [12] results in a noticeably larger energy productivity growth than the CCC scenario, 61% between 2022 and 2050. However, it gives a much smaller overall aggregate impact due to a lower Domar weight than globally, 0.024, the contribution of energy production to aggregate productivity growth is thus $\ln(0.61 + 1) \times 0.024 = 0.011$, that is an increase of 1.1 percentage points over the period from 2022 to 2050 (0.04pp per year on average).

For the scenario based on the CCC scenario, the energy mix leads to even lower productivity growth, about 13%, mostly due to its higher reliance on Carbon Capture and Storage. With the same Domar weight of 0.024, the contribution of energy to aggregate productivity growth is an increase of 0.29pp over 2022–2050, thus 0.01pp per year. The prediction interval overlaps zero, suggesting that these gains are economically small and not statistically different from zero.

Three factors account for the gap between the UK and the World Fast Transition scenarios: higher initial energy prices in the world (24.4 pp), growth in useful energy demand (11.5 pp), and a more favourable solar-to-wind ratio (8.0 pp).

By contrast, the difference between the two UK scenarios is driven entirely by energy-mix composition. Abated electricity generation from fossil fuels or biomass coupled with CCS remains




Figure 3: **Useful-energy pathways and their productivity impact for three net-zero scenarios.** Rows: global Fast Transition based on Way et al. [12] (top; quantities in EJ), UK Fast Transition based on Way et al. [12] (middle; TWh), and UK CCC Sixth Carbon Budget Balanced Net Zero Pathway (bottom; TWh). Columns: **(left)** useful-energy mix by source, 2022–2050, with grid and storage costs (batteries for short/medium duration, hydrogen from solar and wind for long duration) included in the delivered cost of intermittent renewables; **(middle)** unit price of useful energy (£/MWh), constructed as a quantity-weighted index over the mix in the left column using technology-level cost trajectories and technology roll-out from Way *et al.* (2022)—identical across scenarios—with dashed horizontal lines marking the 2022 and 2050 UK Fast Transition prices for reference; **(right)** cumulative TFP growth of the energy sector (pp relative to 2022). Shaded bands are 90% empirical prediction intervals propagating the industry-level forecast error at calibrated $\zeta = 0.13$.

costly, so including these technologies in the mix can reduce productivity growth in the energy sector overall. In the CCC Balanced Net Zero pathway, abated gas and biomass still account for 4.8% of final energy in 2050, whereas they are absent from the pathway in Way et al. [12].

Discussion

In this paper, we develop a parsimonious and flexible framework to estimate the productivity growth impact of the net-zero transition, as implied by bottom-up techno-economic scenarios.

We use only the price and quantity paths reported in techno-economic roadmaps to recover implied industry-level productivity shocks and aggregate them to the economy level without imposing a specific production function or a full general-equilibrium structure. This provides a simple bridge from detailed transition scenarios to aggregate total factor productivity, a macroeconomic object central to long-run income growth but rarely reported directly in transition exercises.

Both steps of the bridge (industry-level mapping and Domar aggregation) are back-tested on historical data, yielding approximately unbiased forecasts whose error distributions supply data-driven prediction intervals. This makes the source of every number much more traceable. In that sense, we view our estimates as more robust than those from more sophisticated integrated assessment models. Our approach therefore provides an empirically disciplined benchmark against which more structural methods should be compared.

The substantive findings overturn a common framing. The net-zero debate is often cast as a contest between techno-optimists who promise large productivity dividends and pessimists who warn of severe economic drag. For the UK, if we believe the CCC techno-economic scenarios, neither position holds when confronted with the data: the net-zero transition is close to productivity-neutral, with cumulative effects by 2050 falling within the measurement error of national productivity statistics.

Yet technology choices still matter. This is particularly true in the energy sector, where extensive electrification and renewable electricity deployment are more likely to produce TFP gains.

The global picture is strikingly different. At the world level, the fast-transition energy scenario generates nearly 6 percentage points of cumulative aggregate productivity gains by 2050, that is, about 0.2 additional percentage points of TFP growth per year. If the UK is representative of the Global North, much of the productivity upside may lie in emerging economies where high initial energy costs, growing demand, and untapped solar potential amplify the returns to transition.

Materials and Methods

Derivation of the reduced-form bridge

The starting point is the industry accounting identity for sector i , total sales equal total expenditures on intermediate inputs X_j and primary factors L_f . Taking time derivatives of the identity $p_i X_i = \sum_j p_j X_{ij} + \sum_f w_f L_{fi}$ gives

$$\hat{p}_i + \hat{X}_i = \sum_j a_{ji}(\hat{p}_j + \hat{X}_{ij}) + \sum_f \tilde{\ell}_{fi}(\hat{w}_f + \hat{L}_{fi}), \quad (5)$$

where p_j and w_f are the price of input j and factor f , $a_{ji} = p_j X_{ij}/(p_i X_i)$ is the cost share of intermediate input j and $\tilde{\ell}_{fi} = w_f L_{fi}/(p_i X_i)$ is the cost share of factor f . Industry TFP growth is the Solow residual, that is, the gap between output growth and cost-share-weighted input growth,

$$\hat{A}_i = \hat{X}_i - \underbrace{\left(\sum_j a_{ji} \hat{X}_{ij} + \sum_f \tilde{\ell}_{fi} \hat{L}_{fi} \right)}_{\equiv \hat{\Sigma}_X}. \quad (6)$$

Equivalently, from the dual side, $\hat{A}_i = -\hat{p}_i + \hat{\Sigma}_p$, where $\hat{\Sigma}_p \equiv \sum_j a_{ji} \hat{p}_j + \sum_f \tilde{\ell}_{fi} \hat{w}_f$ is the cost-share-weighted growth of input prices. The accounting identity of sales growth Eq. 5 can then be written compactly as $\hat{s}_i = \hat{p}_i + \hat{X}_i = \hat{\Sigma}_p + \hat{\Sigma}_X$.

Computing $\hat{A}_i$ directly requires the full vector of input volumes and prices, which transition scenarios do not provide. Our key simplifying assumption replaces this requirement with a single parameter: we assume that cost-share-weighted input-price growth is proportional to nominal sales growth, $\hat{\Sigma}_p = \zeta \hat{s}_i$, with a common proportionality factor ζ across industries and time periods. In other words, we make the key assumption that the relative contribution of input volume growth and input price growth is constant, that is, we assume that there exists

$$\zeta = \frac{\hat{\Sigma}_{p_i}}{\hat{\Sigma}_{p_i} + \hat{\Sigma}_{X_i}}. \quad (7)$$

Since TFP (Eq. 6) can be rewritten as $\hat{A}_i = \hat{X}_i - \hat{\Sigma}_X = \hat{\Sigma}_p - \hat{p}_i$, we can substitute Eq. 7 into the dual expression to yield

$$\hat{A}_i = \zeta \hat{X}_i + (\zeta - 1) \hat{p}_i. \quad (8)$$

Domar aggregation with multiple primary factors

Aggregate TFP growth is obtained by Domar aggregation, $\hat{A} = \sum_i \lambda_i \hat{A}_i$ with $\lambda_i = p_i X_i/(pY)$. The proof extends familiar one-factor results [13] to an economy with labour and capital. Write GDP in value as $pY = \sum_i p_i C_i$ where C_i is final demand in industry i , and define aggregate TFP as the gap between real GDP growth and factor-share-weighted factor growth, $\hat{A} = \sum_i (p_i C_i/pY) \hat{C}_i - \sum_f (w_f L_f/pY) \hat{L}_f$. Substituting the LHS with the formula for Domar aggregation, and then each $\hat{A}_i$ by the primal expression Eq. 6, using the goods-market clearing condition $X_i = \sum_j X_{ji} + C_i$ and rearranging double sums proves the identity. See SI for details.

In our forecasting exercise we hold Domar weights fixed at the forecast origin t . Future expenditure shares are unobserved when the forecast is made, and in the historical backtest, replacing fixed-origin weights with the average of initial and terminal weights does not materially improve forecast accuracy (see SI).

Data construction

We validate the framework on three industry databases: EU KLEMS (baseline), OECD STAN (short-horizon cross-check), and the NBER-CES Manufacturing Industry Database (US manufacturing). Across datasets, annual industry TFP is computed using Törnqvist weights,

$$\hat{A}_{i,t,1} = \Delta \ln X_Q - \bar{\alpha}_L \Delta \ln L_Q - \bar{\alpha}_K \Delta \ln K_Q - \bar{\alpha}_X \Delta \ln Z_Q, \quad (9)$$

where Y_Q, L_Q, K_Q, Z_Q are real gross output, labour, capital services, and intermediate inputs, and $\bar{\alpha}_f = [\alpha_f(t) + \alpha_f(t-1)]/2$ are two-period-average cost shares computed from current-price accounts. Capital expenditure is recovered residually as value added minus labour compensation. Long-horizon TFP growth is accumulated from the resulting TFP levels rather than by summing annual growth rates.

Real output-price growth is measured relative to the country GDP deflator, $\hat{p}_{i,c,t,\tau} = [\ln P_{i,t+\tau}^Y - \ln P_{i,t}^Y] - [\ln P_{c,t+\tau}^{\text{GDP}} - \ln P_{c,t}^{\text{GDP}}]$, so that $\hat{p} > 0$ means the industry's output price rises faster than the aggregate price level. The GDP deflator is the total-economy value-added price index. We exclude Slovakia, Ireland, Greece, and aggregate EU panels due to structural breaks, and restrict to countries with complete first-digit sector coverage, leaving 18 countries with a maximum horizon of 24 years.

Calibration of ζ

For each forecast in country c , industry i between t and $t + \tau$, the out-of-sample error of industry-level TFP growth forecast is $v_{i,c,t,\tau}(\zeta) = [\zeta \hat{X}_{i,c,t,\tau} + (\zeta - 1) \hat{p}_{i,c,t,\tau}] - \hat{A}_{i,c,t,\tau}$. Let $\mathcal{S} = \{(i, c, t, \tau) : \tau \in \mathcal{T}, \text{ data available in KLEMS}\}$ with $\mathcal{T} = \{15, 16, \dots, \bar{\tau}\}$ the set of transition-relevant horizons (where $\bar{\tau}$ is the longest horizon supported by the KLEMS panel for this particular series), and $N = |\mathcal{S}|$. We choose ζ to minimise the pooled mean absolute error

$$\text{MAE}(\zeta) = \frac{1}{N} \sum_{(i,c,t,\tau) \in \mathcal{S}} |v_{i,c,t,\tau}(\zeta)|. \quad (10)$$

The minimiser is $\zeta = 0.13$ when using KLEMS and horizons $\tau \geq 15$ years. Results are similar under alternative loss functions (mean error, median absolute error, root mean squared error, sales-weighted variants) and when estimated with regression models including country, sector, year, and horizon fixed effects (see Supplementary Information for full robustness checks). We obtain similar result for ζ in STAN and NBER-CES, 0.19 in STAN, and 0.13 in NBER-CES, although we prefer to use KLEMS because STAN has a shorter time horizon and NBER-CES is limited to manufacturing.

Construction of prediction intervals

The aggregate forecast error from using predicted rather than true industry TFPs in Domar aggregation can be decomposed exactly as

$$\rho_{c,t,\tau} = \frac{\sum_{j \in \mathcal{I}} \lambda_{j,c,t}}{\sum_i \lambda_{i,c,t}} \varepsilon_{c,t,\tau} + \sum_{j \in \mathcal{I}} \lambda_{j,c,t} v_{j,c,t,\tau}. \quad (11)$$

where ε is the aggregation error, v_j is the industry-level mapping error, and $\mathcal{I}$ is the set of sectors included in the analysis. The reported uncertainty bands are obtained by simulation: (i) draw one aggregation error $\varepsilon^{(b)}$ from the (pooled) empirical distribution of $\varepsilon_{c,t,\tau}$; (ii) for each forecasted industry j , draw one industry-level error $v_j^{(b)}$ from the (pooled) empirical distribution of $v_{j,c,t,\tau}$; (iii) combine as $\rho^{(b)} = \varepsilon^{(b)} + \sum_j \lambda_j v_j^{(b)}$; (iv) repeat 10,000 times. The bands correspond to quantiles of the resulting distribution. Because the KLEMS validation horizon (24 years) is shorter than the longest projection horizon, each error quantile is extrapolated linearly beyond sample support using the slope estimated from the last 10 observed horizons. Note that we assumed that industry-level errors are independent across sectors conditional on the horizon, and independent of aggregation errors.

These intervals quantify the uncertainty introduced by our translation framework, conditional on the scenario paths for prices and quantities. They do not capture the uncertainty in the scenarios themselves.

Author Contributions

E.R. and F.L. designed the study and wrote the manuscript. E.R. collected and analysed the data.

Declaration of Interests

The authors declare no competing interests. F.L. is affiliated scientist with Macrocosm Inc.

Correspondence

¹ Contact: emilien.ravigne@smithschool.ox.ac.uk; francois.lafond@inet.ox.ac.uk.

Data Availability

Replication code and processed datasets are available on a [GitHub repository](#). Raw data are publicly available from EU KLEMS (<https://euklems-intanprod-llce.luiss.it/>), OECD STAN (<https://stats.oecd.org>), NBER-CES (<https://www.nber.org/research/data/>

nber-ces-manufacturing-industry-database), and the Climate Change Committee’s Carbon Budgets (<https://www.theccc.org.uk/publicationtype/report/carbon-budget/>).

Acknowledgments

We gratefully acknowledge funding from UK Research and Innovation through Economic and Social Research Council grant PRINZ (ES/W010356/1), from Baillie Gifford, and from the Institute for New Economic Thinking at the Oxford Martin School, as well as comments from our IARIW discussant (Alicia Rambaldi) and audiences at Oxford, the ECB Working Group on Economic Modelling, CSH Vienna, EIEE-CMCC, IEW, EAERE, Collegio Carlo Alberto and the Schumpeter conference.

We used modern AI assistants for manuscript preparation.

References

- [1] Sam Fankhauser, Stephen M. Smith, Myles Allen, Kaya Axelsson, Thomas Hale, Cameron Hepburn, J. Michael Kendall, Radhika Khosla, Javier Lezaun, Eli Mitchell-Larson, Michael Obersteiner, Lavanya Rajamani, Rosalind Rickaby, Nathalie Seddon, and Thom Wetzer. The meaning of net zero and how to get it right. *Nature Climate Change*, 12(1):15–21, January 2022. ISSN 1758-6798. doi: 10.1038/s41558-021-01245-w.
- [2] Keywan Riahi, Roberto Schaeffer, Julián Arango, Katherine Calvin, Céline Guivarch, Tomoko Hasegawa, Kejun Jiang, Elmar Kriegler, Robin Matthews, Glen P. Peters, Anjali Rao, Simon Robertson, Anne M. Sebbit, Julia Steinberger, Massimo Tavoni, and Detlef P. van Vuuren. Mitigation pathways compatible with long-term goals. In P. R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, and J. Malley, editors, *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK and New York, NY, USA, 2022. doi: 10.1017/9781009157926.005. URL <https://www.ipcc.ch/report/ar6/wg3/chapter/chapter-3/>.
- [3] International Energy Agency. Net zero by 2050: A roadmap for the global energy sector. Technical report, International Energy Agency, Paris, May 2021. URL <https://www.iea.org/reports/net-zero-by-2050>.
- [4] CCC. The Seventh Carbon Budget: advice for the government. Technical report, Committee on Climate Change, February 2025. URL <https://www.theccc.org.uk/wp-content/uploads/2025/02/The-Sevenh-Carbon-Budget.pdf>.
- [5] Gert Bijnens, Sofia Anyfantaki, Andrea Colciago, Jan De Mulder, Elisabeth Falck, Vincent Labhard, Paloma Lopez-Garcia, Nuno Lourenço, Jaanika Meriküll, Miles Parker, et al. The impact of climate change and policies on productivity. *ECB Occasional Paper*, 2024(340), 2024.
- [6] Nicola Benatti, Martin Groiss, Petra Kelly, and Paloma Lopez-Garcia. Environmental regulation and productivity growth in the euro area: Testing the Porter hypothesis. *Journal of Environmental Economics and Management*, 126:102995, July 2024. ISSN 0095-0696. doi: 10.1016/j.jeem.2024.102995.

- [7] Jean-Francois Mercure, Hector Pollitt, Frank W Geels, and Dimitri Zenghelis. The effects of low-carbon transitions on labour productivity: analysing uk electricity, heat, and mobility with a techno-economic simulation model. *Climate Policy*, 26(4):630–647, 2026.
- [8] S. Batten and S. Millard. Productivity implications of the move to net zero. Working Paper 065, The Productivity Institute, 2026. URL <https://www.productivity.ac.uk/wp-content/uploads/2026/02/WP065-FINAL.pdf>.
- [9] Stephane Hallegatte, Floren McIsaac, Hasan Dudu, Charl Jooste, Camilla Knudsen, and Hans Beck. The macroeconomic implications of a transition to zero net emissions: A modeling framework. *World Bank Policy Research Working paper*, 10367, 2023.
- [10] Costas Arkolakis and Conor Walsh. The economic impacts of clean energy. *Brookings Papers on Economic Activity*, 2024(2):161–212, 2024.
- [11] CCC. The Sixth Carbon Budget: The UK’s path to Net Zero. Technical report, Committee on Climate Change, December 2020. URL <https://www.theccc.org.uk/wp-content/uploads/2020/12/The-Sixth-Carbon-Budget-The-UKs-path-to-Net-Zero.pdf>.
- [12] Rupert Way, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer. Empirically grounded technology forecasts and the energy transition. *Joule*, 6(9):2057–2082, September 2022. ISSN 2542-4351. doi: 10.1016/j.joule.2022.08.009.
- [13] James McNerney, Charles Savoie, Francesco Caravelli, Vasco M. Carvalho, and J. Doyne Farmer. How production networks amplify economic growth. *Proceedings of the National Academy of Sciences*, 119(1):e2106031118, January 2022. doi: 10.1073/pnas.2106031118.
- [14] Bert M Balk. *Productivity*. Springer, 2021.
- [15] Jing Lei, Max G’Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.
- [16] Yun Shin Lee and Stefan Scholtes. Empirical prediction intervals revisited. *International Journal of Forecasting*, 30(2):217–234, 2014.
- [17] J Doyne Farmer and Francois Lafond. How predictable is technological progress? *Research policy*, 45(3):647–665, 2016.
- [18] IEA. Energy Technology Perspectives 2023. Technical report, International Energy Agency, January 2023.
- [19] MPP. Making Net Zero Concrete and Cement Possible. Technical report, Mission Possible Partnership, 2023.
- [20] McKinsey. Green cement and the net-zero transition. Technical report, McKinsey & Company, January 2022.
- [21] HM Treasury. The green book: Central government guidance on appraisal and evaluation. Technical report, HM Treasury, 2022. URL <https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government>. Accessed: 2026-04-27.
- [22] Ian Goldin, Pantelis Koutroumpis, François Lafond, and Julian Winkler. Why is productivity slowing down? *Journal of Economic Literature*, 62(1):196–268, 2024. doi: 10.1257/jel.20221543.
- [23] Nicholas Oulton. Must the growth rate decline? baumol’s unbalanced growth revisited. *Oxford Economic Papers*, 53(4):605–627, 2001.
- [24] J Stan Metcalfe, John Foster, and Ronnie Ramlogan. Adaptive economic growth. *Cambridge journal of Economics*, 30(1):7–32, 2006.
- [25] François Lafond, Diana Greenwald, and J Doyne Farmer. Can stimulating demand drive costs down? world war ii as a natural experiment. *The Journal of Economic History*, 82(3):727–764, 2022.

- [26] Bert M Balk. *Price and quantity index numbers: models for measuring aggregate change and difference*. Cambridge University Press, 2012.
- [27] Vasco Carvalho and Xavier Gabaix. The great diversification and its undoing. *American Economic Review*, 103(5):1697–1727, 2013.
- [28] Andrew T Foerster, Pierre-Daniel G Sarte, and Mark W Watson. Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy*, 119(1):1–38, 2011.
- [29] CCC. Methodology Report for The Sixth Carbon Budget. Technical report, Committee on Climate Change, December 2020.
- [30] Climate Change Committee. Methodology report: Methodology behind the ccc’s carbon budget advice for the uk, northern ireland, wales, and scotland, published in 2025. Technical report, Climate Change Committee, May 2025. URL <https://www.theccc.org.uk/publications/>. Accessed: 2026-04-27.
- [31] Lion Hirth. The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy economics*, 38:218–236, 2013.
- [32] François Lafond, Aimee Gotway Bailey, Jan David Bakker, Dylan Rebois, Rubina Zadourian, Patrick McSharry, and J Doyne Farmer. How well do experience curves predict technological progress? a method for making distributional forecasts. *Technological Forecasting and Social Change*, 128:104–117, 2018.
- [33] Carbon Brief. In-depth: The whole system costs of renewables. <https://www.carbonbrief.org/in-depth-whole-system-costs-renewables/>, February 2017.
- [34] Department for Business, Energy & Industrial Strategy. Energy consumption in the uk 2022. Statistical report, UK Government, 2022. URL <https://www.gov.uk/government/statistics/energy-consumption-in-the-uk-2022>. Accessed: 2026-04-16.

SUPPLEMENTARY INFORMATION

This S.I provides all the details and derivations. It is organised as

I	Theoretical framework	S2
1	Domar aggregation	S2
2	Industry-level TFP	S4
A	Estimating industry-level productivity growth from prices and volume scenarios	S4
B	There exists a universal ζ that recovers <i>aggregate</i> productivity	S5
C	Behavioral and economic interpretation of the assumption $\zeta_{it} = \zeta$	S6
D	Price index and goods substitution	S6
E	Real price growth	S8
3	Partial aggregation of the economy	S10
4	Final versus useful energy	S11
5	Empirical validation of forecasts	S13
A	Data construction	S13
B	Validation of fixed-weight Domar aggregation	S14
C	Validation of the industry-level mapping from prices and output to industry TFP growth . .	S15
D	Robustness to the calibration objective	S17
E	Econometric validation of ζ in KLEMS	S18
F	Combining industry-level and aggregation uncertainty	S19
G	Separating Sectoral Calibration from Domar Aggregation	S19
II	Applications	S22
6	Climate Change Committee scenarios	S23
A	Framework	S23
B	Results	S26
C	Accounting conventions	S26
D	Robustness	S27
E	Propagation of validation errors in the CCC exercise	S27
7	Energy	S29
A	Input prices and cost assumptions	S29
B	Scenario Quantities	S31
C	Main useful-energy results	S32
D	Counterfactual baselines	S32
E	Robustness to calibration assumptions	S34

Part I

Theoretical framework

1 Domar aggregation

McNerney et al. [13] derive Domar aggregation from a basic accounting framework. However, their derivations consider only a single factor, labour, so their equations typically express real quantities by deflating using the wage rate. To ensure that we interpret and measure TFP shocks properly, we extend the framework to include capital. We also provide a much more straightforward proof.

Starting from the industry-level accounting identity,

$$\underbrace{\sum_{j=1}^N X_{ji} p_i + \sum_{l=1}^D C_{li} p_i}_{\text{Total Sales, } p_i X_i} = \underbrace{\sum_{j=1}^N X_{ij} p_j + \sum_{f=1}^F L_{fi} w_f}_{\text{Total expenses}},$$

we can take time derivatives and rearrange to get

$$\hat{p}_i + \hat{X}_i = \sum_{j=1}^N a_{ji} (\hat{p}_j + \hat{X}_{ij}) + \sum_{f=1}^F \tilde{l}_{fi} (\hat{L}_{fi} + \hat{w}_f),$$

where the a_{ji} and $\tilde{l}_{fi}$ are the shares of intermediate input j and factor f in i 's expenses. Rearranging, we have the definition of industry-level TFP as

$$\hat{A}_i \equiv \hat{X}_i - \underbrace{\sum_{j=1}^N a_{ji} \hat{X}_{ij}}_{\text{TFP, primal approach}} - \underbrace{\sum_{f=1}^F \tilde{l}_{fi} \hat{L}_{fi}}_{\text{TFP, dual approach}} = - \left(\hat{p}_i - \sum_{j=1}^N a_{ji} \hat{p}_j - \sum_{f=1}^F \tilde{l}_{fi} \hat{w}_f \right). \quad (12)$$

GDP (in value) is the sum of final consumption

$$\text{GDP} = pY = \sum_{i=1}^N p_i C_i.$$

Taking time derivatives,

$$\hat{p} + \hat{Y} = \sum_{i=1}^N \frac{p_i C_i}{pY} \hat{C}_i + \sum_{i=1}^N \frac{p_i C_i}{pY} \hat{p}_i = \sum_{f=1}^F \frac{w_f L_f}{pY} \hat{L}_f + \sum_{f=1}^F \frac{w_f L_f}{pY} \hat{w}_f,$$

so that assuming a price index

$$\hat{p} = \sum_{i=1}^N \frac{p_i C_i}{pY} \hat{p}_i,$$

we have real GDP growth as

$$\hat{Y} = \sum_{i=1}^N \frac{p_i C_i}{pY} \hat{C}_i. \quad (13)$$

Now we *define* aggregate TFP growth as the growth of (volume) GDP minus the factor-share weighted growth

of the volume of factors

$$\hat{A} \equiv \underbrace{\sum_{i=1}^N \frac{p_i C_i}{pY} \hat{C}_i}_{\text{growth of real GDP}} - \underbrace{\sum_{f=1}^F \frac{w_f L_f}{pY} \hat{L}_f}_{\text{growth of volume of factors}}. \quad (14)$$

We want to show that Domar aggregation is valid, that is

$$\hat{A} = \sum_{i=1}^N \lambda_i \hat{A}_i, \quad (15)$$

where the “Domar” weight $\lambda_i \equiv p_i X_i / pY$.

Proof

We want to show that the definition of aggregate TFP, Eq. 14, is equal to Eq. 15. We show this using the primal side of TFP. Substituting the LHS of Eq. 12 into Eq. 15 and replacing λ_i by its definition, we need to prove

$$\hat{A} = \sum_{i=1}^N \frac{p_i X_i}{pY} \left(\hat{X}_i - \sum_{j=1}^N \frac{p_j X_{ij}}{p_i X_i} \hat{X}_{ij} - \sum_{f=1}^F \frac{w_f L_{fi}}{p_i X_i} \hat{L}_{fi} \right) = \sum_{j=1}^N \frac{p_j C_j}{pY} \hat{C}_j - \sum_{f=1}^F \frac{w_f L_f}{pY} \hat{L}_f. \quad (16)$$

Distributing the LHS into three sums and simplifying all the $p_i X_i / p_i X_i$,

$$\hat{A} = \underbrace{\sum_{i=1}^N \frac{p_i X_i}{pY} \hat{X}_i}_{(1)} - \underbrace{\sum_{i=1}^N \sum_{j=1}^N \frac{p_j X_{ij}}{pY} \hat{X}_{ij}}_{(2)} - \underbrace{\sum_{i=1}^N \sum_{f=1}^F \frac{w_f L_{fi}}{pY} \hat{L}_{fi}}_{(3)}. \quad (17)$$

Consider the third term, and permute the sums to get

$$(3) = \sum_{f=1}^F \frac{w_f}{pY} \sum_{i=1}^N L_{fi} \hat{L}_{fi} = \sum_{f=1}^F \frac{w_f L_f}{pY} \hat{L}_f, \quad (18)$$

where we have used $L_f \hat{L}_f = \sum_i L_{fi} \hat{L}_{fi}$. Similarly for the second term, permuting the two sums and using the fact that since $X_j = \sum_i X_{ij} + C_j$, $\sum_i X_{ij} \hat{X}_{ij} = X_j (\hat{X}_j - (C_j / X_j) \hat{C}_j)$, it becomes

$$(2) = \sum_{j=1}^N \frac{p_j}{pY} \sum_{i=1}^N X_{ij} \hat{X}_{ij} = \sum_{j=1}^N \frac{p_j X_j}{pY} \hat{X}_j - \sum_{j=1}^N \frac{p_j C_j}{pY} \hat{C}_j. \quad (19)$$

Substituting Eqs 18 and 19 back in Eq. 17, the first term of Eq. 19 cancels out the first term in Eq. 17, so that we are left with the RHS of Eq. 16. #

Oulton [23, p.615] also derives Domar aggregation, but starting from production functions and assuming competition, rather than simply defining TFP as a Solow residual, as we do here, in the spirit of Balk [14] and McNerney et al. [13].

2 Industry-level TFP

Here we elaborate on the empirical and theoretical content of our assumption relating industry-level prices and volume growth to productivity growth, Eq. 24.

A Estimating industry-level productivity growth from prices and volume scenarios

Since Domar aggregation tells us how to aggregate industry-level productivity shocks, all we have to do now is to estimate these industry-level productivity shocks. However, estimating productivity growth from Eq. 12 is a daunting task, as it requires estimating the change in unit requirement for all inputs. Our key contribution here is to introduce a simple, transparent, and testable assumption that allows us to circumvent this problem.

First, note that net-zero scenarios typically include pathways for the prices and/or output of the industry of interest, $\hat{p}_i$ and $\hat{X}_i$. So we are looking for a way to predict TFP in the future using only future values of x_i and p_i .

Let us start by noting that TFP growth can be written equivalently in its primal or dual form,

$$\hat{A}_i \equiv \hat{X}_i - \sum_{j=1}^N a_{ji} \hat{X}_{ij} - \sum_{f=1}^F \tilde{l}_{fi} \hat{L}_{fi} = -\left(\hat{p}_i - \sum_{j=1}^N a_{ji} \hat{p}_j - \sum_{f=1}^F \tilde{l}_{fi} \hat{w}_f \right). \quad (20)$$

In plain english, total factor productivity growth can be seen as either the growth of volume, netted out of the volume growth of inputs, or as the decline in price, netted out of the decline in the prices of inputs. Introducing the notation

$$\hat{\Sigma}_{X_i} \equiv \sum_{j=1}^N a_{ji} \hat{X}_{ij} + \sum_{f=1}^F \tilde{l}_{fi} \hat{L}_{fi}$$

for the growth of the index of inputs of i , and

$$\hat{\Sigma}_{p_i} \equiv \sum_{j=1}^N a_{ji} \hat{p}_j + \sum_{f=1}^F \tilde{l}_{fi} \hat{w}_f$$

for the growth of the index of prices of the inputs of i , the primal-dual identity of TFP, Eq. 20, can be rewritten as

$$\hat{A}_i = \hat{X}_i - \hat{\Sigma}_{X_i} = \hat{\Sigma}_{p_i} - \hat{p}_i. \quad (21)$$

Rearranging, this means that the growth of (nominal) sales can be written as

$$\hat{s}_i = \hat{X}_i + \hat{p}_i = \hat{\Sigma}_{X_i} + \hat{\Sigma}_{p_i}. \quad (22)$$

In other words, the growth rate of nominal sales is equal to the growth rate of the cost of inputs, plus the growth rate of the volume of inputs. We make the key assumption that the relative contribution of these two terms is constant, that is, we assume that there exists

$$\zeta = \frac{\hat{\Sigma}_{p_i}}{\hat{\Sigma}_{p_i} + \hat{\Sigma}_{X_i}}. \quad (23)$$

If ζ was indexed by industry and time, Eq. 23 would just be a definition. However, we assume that ζ is constant over time, country and industry. We could have used industry-specific and/or country-specific ζ . We prefer to use a universal value, because while it is possible that industry-specific values of ζ would lead to more accurate predictions, in practice this also introduces more noise so we face a classic bias-variance trade-off. This is also simpler and thus more transparent.

Substituting assumption 23 into Eq. 21, and rearranging using Eq. 22 leads

$$\hat{A}_i = \zeta \hat{X}_i + (\zeta - 1) \hat{p}_i. \quad (24)$$

Eq. 24 is extremely useful – it allows us to determine the growth rate of TFP of industry i by knowing only its price and output growth – we do not need the price or volume growth of its inputs. Eq. 24 means that productivity growth translates into output growth and price decrease (if $0 < \zeta_i < 1$) in fixed proportions.

Eq. 24 is a purely statistical relation, with little motivation on theoretical grounds; in other words, it is a reduced form relation which is not immune to the Lucas critique. However, this is not an issue here - we do not need to assume that this relationship is causal, and we do not proceed to any comparative static; we only need a statistical relation between price/output growth and productivity growth, allowing us to infer the productivity growth assumptions that are implicit in net-zero scenarios, which only provide future prices and quantities. Appendix 2 elaborates on this point further.

Appendix 5 provides a thorough evaluation of the validity of Eq. 24 in past data using backtesting methods, for a universal value of ζ . That is, for various databases, we put ourselves in the past and assume that we know future values of prices and output, and try to predict productivity, just as we will do in the next section with net-zero scenarios. Crucially, we show that we can find ζ such that Eq. 24 leads to almost unbiased forecasts. Furthermore, we use the forecast errors to characterise uncertainty in the next section. In other words, we will make predictions of productivity conditional on prices and quantities, and we will assume that our errors for these predictions are similar to the errors we would have made in the past if we had used this method.

B There exists a universal ζ that recovers aggregate productivity

First of all, it is helpful to note that there exists ζ^* such that using Eq. 24 to predict industry-level TFPs will still lead to a correct prediction for *aggregate* TFP. To show this, let us take Domar aggregation as accounting truth,

$$\hat{A} = \sum_{i=1}^N \lambda_i \hat{A}_i, \quad (25)$$

and define

$$\hat{\hat{A}}_i \equiv (\zeta \hat{X}_i + (\zeta - 1) \hat{p}_i). \quad (26)$$

By definition of ζ_i ,

$$\hat{A}_i \equiv (\zeta_i \hat{X}_i + (\zeta_i - 1) \hat{p}_i). \quad (27)$$

To find the value of ζ such that using the industry TFP approximation Eq. 26 instead of the true industry Eq. 27 into the Domar aggregation formula Eq. 25 gives the same aggregate productivity, we write

$$\begin{aligned} \sum_{i=1}^N \lambda_i \hat{A}_i &= \sum_{i=1}^N \lambda_i \hat{\hat{A}}_i, \\ \sum_{i=1}^N \lambda_i (\zeta_i \hat{X}_i + (\zeta_i - 1) \hat{p}_i) &= \sum_{i=1}^N \lambda_i (\zeta \hat{X}_i + (\zeta - 1) \hat{p}_i), \\ \sum_{i=1}^N \lambda_i \zeta_i (\hat{X}_i + \hat{p}_i) &= \sum_{i=1}^N \lambda_i \zeta (\hat{X}_i + \hat{p}_i), \end{aligned} \quad (28)$$

Solving for ζ , we have

$$\zeta = \frac{\sum_{i=1}^N \lambda_i \zeta_i (\hat{X}_i + \hat{p}_i)}{\sum_{i=1}^N \lambda_i (\hat{X}_i + \hat{p}_i)}. \quad (29)$$

Thus, there exists a value of ζ , defined as a weighted average of ζ_i , that leads to the same value of aggregate TFP as Domar aggregation of the *true* TFPs. This derivation is helpful in that it clarifies that there is no fundamental accounting issues introduced by the use of Eq. 24 in predicting industry-level TFP, in so far as

one can find the correct value ζ^* .

However, we still face three issues. First, recalling that the ζ_i s are defined as ratios of growth rates, they are very volatile, and we have found empirically that it is not very helpful to try to estimate universal values of ζ using Eq. 29. Second, in practice, the “correct” values ζ_i^* are *period-specific*. To be useful, we need a value of ζ in the future. Third, the aggregate-recovering value in Eq. 29 is implicitly a Domar-weighted object: it can recover historical aggregate productivity even if it does not provide the best mapping from sectoral price and quantity paths into sectoral productivity. We return to this issue below when explaining why our baseline calibration targets the sector-level bridge rather than the full Domar-weighted aggregate pipeline (SI 5G).

As a result, we will estimate ζ so as to make our *industry-level* predictions of productivity as accurate as possible. This allows us to interpret $\zeta \approx \zeta_i$, and discuss the impact of the transition on individual industries rather than simply on the aggregate. In the next section, we discuss this assumption in more details.

C Behavioral and economic interpretation of the assumption $\zeta_{it} = \zeta$

One way to think about this assumption behaviourally is that Eq. 24 simply describes the extent to which productivity growth is translated into either volume growth ($\zeta \approx 1$) or price declines ($\zeta \approx 0$), although we do not restrict $0 < \zeta < 1$. Eq. 24 also shows that, assuming $\zeta > 0$, productivity growth depends positively on output growth, as has been well documented for labour productivity under the name of the Fabricant-Kaldor-Verdoorn law [24], or in the “learning curve” literature when considering the growth of *cumulative* output (see Lafond et al. [25] for a discussion of the differences between the two).

To get further insights into the mechanistic implications of this assumption, we can transform Eq. 22 into

$$\hat{X}_i + \hat{p}_i = \frac{\hat{\Sigma}_p}{\zeta} = \frac{\hat{\Sigma}_X}{1 - \zeta}. \quad (30)$$

If the firm increases its nominal sales by 10%, Eq. 30 implies that the unit cost of inputs grew by $\hat{\Sigma}_p = 10\% \times \zeta$, *no matter whether the sales of the firm increased because of higher volumes or because of higher prices*. While we intuitively think that firms pass down higher input prices as higher output prices, or higher volume inputs purchases as higher volume output, these relationships are broken down here.

Another way to highlight an issue is this. Eq. 30 also implies

$$\hat{\Sigma}_X = \frac{1 - \zeta}{\zeta} \hat{\Sigma}_p, \quad (31)$$

which can be read as a price elasticity of input demand relation, making the assumption of a constant ζ feel somewhat more familiar. However, generally (but not necessarily) we expect $0 < \zeta < 1$, which would imply a *positive* price elasticity. This makes it very clear that ζ does not characterise a behavioural relation, but simply the fact that in the data we generally expect some inflation, and some growth in volume, so that an observed relation like Eq. 31 makes perfect sense.

Again, and to conclude, we regard this relation as a useful reduced-form regularity, and we do not think of ζ as a “deep” parameter. How an industry translates its productivity gains into higher production or lower prices is likely to depend on market characteristics, especially competition, the elasticity of demand, and the rate of innovation. However, we think these characteristics are unlikely to drastically change in the medium term despite the net-zero transition, so that the reduced form relation should remain a useful regularity.

D Price index and goods substitution

Typically, net-zero scenarios describe a process whereby, at the industry level, a green technology replaces a dirty incumbent. Thus, we need to take into account the productivity impact of both the rise of the clean technology *and the decline of the dirty technology*. One approach could be to estimate Eq. 24 on the clean and dirty subsectors separately. However, this raises an issue with using Domar aggregation (Eq. 15), as the Domar weights of the

clean and dirty subindustries are, by definition of a “clean transition”, changing substantially, so we would need to update them dynamically.

Instead, we propose to consider the clean and dirty alternatives together within each industry. That is, we construct an aggregate price and quantity index for the industry (say “electricity” or “steel”), and use the price and quantity index directly in Eq. 24 to determine industry-level productivity growth.

This raises an important question: what is the best index number for doing this? We have found that classic index numbers (Paasche, Laspeyres, Fisher, etc..) can give pathological results, in the sense that substituting a cheap dirty technology by an expensive clean technology can give price declines and thus positive productivity growth. The intuition for this is simple: if both prices are non-increasing, a classic price index will be non-increasing; however, in the case of the net-zero transition, we are sometimes substituting a cheap product by an expensive one, so we would expect the industry-level price to increase.

Fortunately, here, we can make the assumption that each sector has a single, homogeneous output, even though it may arise from different production techniques. For example, the electricity sector produces MWh of electricity, whether they come from fossil fuels or renewables. Similarly for the steel sector, a ton of clean steel from electric furnaces replaces a ton of dirty steel produced using coal.

Therefore, we use what is known as the “unit value” index [26]: we express clean and dirty production in the same units, and define aggregate volume output as the simple sum. The price index is then the average of output prices weighted by their volumes.

To see why we prefer to assume homogeneous output than to use index numbers, consider a simple but pathological example (Table 3).

Table 3: Simple example of a two-good economy

	Period 1	Period 2
<i>Quantities</i>		
Dirty	90	50
Clean	10	50
Total	100	100
<i>Unit cost</i>		
Dirty	600	600
Clean	1500	1200
<i>Sales</i>		
Dirty	54000	30000
Clean	15000	60000
Total	69000	90000

Our economy features a transition from a dominant dirty sector to a prominent clean sector, and the clean sector is initially more expensive but is getting cheaper. Because both prices are non-increasing, typical axiomatic requirements for classic prices indices (Laspeyre, Paasche, etc.) imply that they are non-increasing, see Table 4. However, this is an issue here because the price of an “average unit” is, indeed, increasing.

Table 4: Price and quantity indices

	Laspeyres	Paasche	Fischer	Törnqvist	Unit Value
Price index	0.96	0.86	0.91	0.91	1.30
Quantity index	1.52	1.36	1.44	1.47	1
Sales index	1.30	1.30	1.30	1.33	1.30

Note: The indices are computed using the numerical values in Table 3, and the definitions of the price indices can be found in Balk [26]. Sales index is the product of the price and the quantity indices for Fischer, Törnqvist and the Unit Value Index. For Paasche and Laspeyres, the sales index is the product of the Paasche price index and the Laspeyres quantity index and conversely.

The Paasche price index for instance would be $(50 \times 600 + 50 \times 1200) / (50 \times 600 + 50 \times 1500) = 0.86$, indicating a decrease of roughly 15% of the price, even though we have replaced 40% of the conventional production by a more expensive clean good.

This somewhat pathological outcome arises because price indices are designed to aggregate goods that are of different nature. When goods are homogeneous, we can use the “unit value” index [26]: the price index is the average price paid for one unit of good, and the quantity index is the total number of units.

To sum up, because we need to take into account both the rise of clean goods and the decline of dirty goods, we need to work at the level of a homogeneous sector where a clean good replaces a dirty good, and within this, use unit value price and quantity indices.

E Real price growth

Our framework infers productivity from sectoral price changes. A key measurement issue is whether these price changes should be measured in nominal or real terms. The object of interest is technological change, not variation in input prices. For example, an increase in interest rates raises the cost of capital and can increase prices across many sectors even if production technologies are unchanged. Or a sudden doubling of all prices would not influence production. Interpreting such nominal price movements as evidence of lower productivity would therefore be misleading.

To isolate the technological component, we work with real price changes, defined as nominal price growth net of factor-price inflation. Following McNerney et al. [13], but extending the concept to multiple factors, we define sector-specific factor-price inflation as

$$\hat{r}_i \equiv \sum_f l_{fi} \hat{w}_f, \quad (32)$$

where $l_{fi} = \frac{L_{fi} w_f}{VA_i}$ is the share of factor f in sector i 's value added, and $\sum_f l_{fi} = 1$. The relevant price object is then $\hat{p}_i - \hat{r}_i$: output price growth relative to the growth of factor costs.

This adjustment integrates naturally into the dual definition of productivity. Starting from the standard dual expression and subtracting sector-specific factor-price inflation, one obtains (see proof below)

$$\hat{A}_i = - \left[(\hat{p}_i - \hat{r}_i) - \sum_j a_{ji} (\hat{p}_j - \hat{r}_i) \right]. \quad (33)$$

Thus, the dual formulation can be written in terms of real rather than nominal prices: productivity growth depends on output prices relative to input costs, net of upstream real price effects.

Combining Eq. 33 with the primal representation yields the real-price analogue of our main reduced-form

relationship. Let

$$\hat{\Sigma}_{p_i, \text{real}} \equiv \sum_j a_{ji} (\hat{p}_j - \hat{r}_i).$$

Then

$$\hat{A}_i = \hat{X}_i - \hat{\Sigma}_{X_i} = \hat{\Sigma}_{p_i, \text{real}} - (\hat{p}_i - \hat{r}_i),$$

and hence real sales growth satisfies

$$\hat{s}_i = \hat{X}_i + (\hat{p}_i - \hat{r}_i) = \hat{\Sigma}_{X_i} + \hat{\Sigma}_{p_i, \text{real}}.$$

Defining

$$\zeta \equiv \frac{\hat{\Sigma}_{p_i, \text{real}}}{\hat{\Sigma}_{p_i, \text{real}} + \hat{\Sigma}_{X_i}},$$

we obtain

$$\hat{A}_i = \zeta \hat{X}_i + (\zeta - 1)(\hat{p}_i - \hat{r}_i). \quad (34)$$

Eq. 34 shows that the mapping from prices to productivity is unchanged once nominal price growth is replaced by real price growth. The economic interpretation is immediate. If factor prices rise uniformly, then output prices may also rise even though production efficiency is unchanged. In that case, $\hat{p}_i - \hat{r}_i$ remains small, so the framework does not spuriously infer a productivity loss. By contrast, using nominal prices would attribute such input-cost shocks to productivity.

For consistency, we therefore use real price changes in both the calibration of ζ and the empirical implementation of the framework. This is also aligned with the CCC cost series, which are reported in constant prices. Using nominal price growth would mix real cost wedges with nominal input-price movements and would therefore distort the inferred relationship between prices and productivity.

Derivation of Eq. 33 Starting from the standard dual definition of TFP growth,

$$\hat{p}_i = -\hat{A}_i + \sum_j a_{ji} \hat{p}_j + \sum_f \tilde{l}_{fi} \hat{w}_f,$$

subtract sector-specific factor-price inflation $\hat{r}_i = \sum_f l_{fi} \hat{w}_f$ from both sides:

$$\hat{p}_i - \hat{r}_i = -\hat{A}_i + \sum_j a_{ji} \hat{p}_j + \sum_f (\tilde{l}_{fi} - l_{fi}) \hat{w}_f.$$

Using $\tilde{l}_{fi} = \frac{L_{fi} w_f}{p_i X_i}$, $l_{fi} = \frac{L_{fi} w_f}{\sqrt{A_i}}$, and $a_{ji} = \frac{X_{ij} p_j}{p_i X_i}$, one can rearrange the final term to show that

$$\sum_f (\tilde{l}_{fi} - l_{fi}) \hat{w}_f = - \sum_j a_{ji} \hat{r}_i.$$

Substituting this into the previous expression gives

$$\hat{p}_i - \hat{r}_i = -\hat{A}_i + \sum_j a_{ji} (\hat{p}_j - \hat{r}_i),$$

which rearranges to Eq. 33. #

3 Partial aggregation of the economy

An important issue is that our method makes it possible to predict industry-level productivity based on prices and quantities alone, without the need for technical coefficients. However, it does not solve the issue of coverage - to estimate the impact of the net zero transition on the aggregate economy, we still need to estimate the impact on each affected industry. In other words, let us suppose that the net zero transition impacts productivity in 3 sectors $\{A, B, C\}$. Aggregate productivity is

$$\hat{A} = \sum_{i \in \{A, B, C\}} \lambda_i \hat{A}_i,$$

and our method offers an approximation via

$$\hat{A} \approx \sum_{i \in \{A, B, C\}} \lambda_i \tilde{A}_i.$$

We still must consider all sectors where productivity has changed, in the sense that

$$\hat{A} \approx \sum_{i \in \{A, B\}} \lambda_i \tilde{A}_i,$$

for instance, gives an incomplete estimate.

Our historical validation uses economy-wide industry data, so Domar aggregation applies over the complete economy. In scenario applications, however, there can be a concern that we do not have complete coverage of all the sectors where productivity has changed.

While we recognise that this issue likely biases our results, we think our estimates are still useful for two reasons. First, in the CCC scenarios, the sectors represented cover a large part of the economy, and are in fact designed to represent all the sectors affected by the transition, so the issue of limited coverage should be relatively limited. In the case of the application to electricity, we deliberately take a specific industry, so it is clear that our estimates concern only this part of the economy.

Second, there could be a concern that generically, productivity improvements in one part of the economy cause productivity changes in another part of the economy. In practice, cross-sectional in TFP shocks tend to be relatively limited (and much more limited than correlations in prices or output). For instance, [27] report an average correlation in TFP shocks of only 2.3% among 77 US private industries, and taking TFP correlations into account helps only modestly to understand aggregate fluctuations. In addition, existing measured cross-sectional correlations in TFP shocks might be due to common factors rather than bilateral cross-industry dependence, as suggested by the modelling choices in [28], for instance. That said, it remains plausible that innovation in one industry impacts innovation in another, so this could constitute a limitation of our approach.

4 Final versus useful energy

A transition from fossil fuels to electricity changes *final* energy use because conversion losses differ across technologies, even when delivered *useful* energy is unchanged. Since our framework must choose one output concept for the energy sector, we clarify here why this choice does not affect aggregate productivity in the full-economy accounting, and why using useful energy is the safer empirical option.

We illustrate the point in a two-sector economy with an energy sector E and a downstream sector C . Suppose the economy switches from fossil energy F to electricity G , with

$$X_G = xX_F, \quad 0 < x < 1,$$

so that both inputs provide the same useful energy. The case $x < 1$ captures higher end-use efficiency under electrification. For simplicity, assume the unit price of electricity satisfies

$$p_G = \frac{1}{x}p_F,$$

so that the expenditure of sector C on energy is unchanged and the price of 1 useful MWh is constant across energies.

Final-energy accounting. If energy output is measured in final-energy units, the energy sector's quantity falls by $\ln x$ and its price rises by $-\ln x$:

$$\hat{X}_E = \ln x, \quad \hat{p}_E = -\ln x.$$

Its productivity growth is therefore negative,

$$\hat{A}_E = \ln x.$$

For the downstream sector, physical output is unchanged, $\hat{X}_C = 0$, but the quantity of energy input falls while its price rises. Let w_{EC} denote the energy input share in sector C . Then

$$\hat{A}_C = -w_{EC} \ln x.$$

Aggregate productivity growth is the Domar-weighted sum

$$\hat{A} = \lambda_E \hat{A}_E + \lambda_C \hat{A}_C = \lambda_E \ln x - \lambda_C w_{EC} \ln x.$$

Using the identity that the sales of E are the electricity expenditures of C , $p_C X_C w_{EC} = p_E X_E$, which implies $\lambda_C w_{EC} = \lambda_E$, we obtain

$$\hat{A} = 0.$$

Useful-energy accounting. If instead the energy sector is defined as producing useful energy, neither the quantity nor the price of energy services changes. The energy sector supplies the same useful-energy bundle before and after the switch, so

$$\hat{X}_E = \hat{p}_E = \hat{A}_E = 0.$$

Sector C also receives the same quantity of useful energy at the same price, implying

$$\hat{A}_C = 0.$$

Hence aggregate productivity growth is again

$$\hat{A} = 0.$$

Useful-energy accounting is safer only conditionally on non overlapping sectors. The accounting is exact when the energy sector is defined in service units, but it can overstate aggregate productivity gains if downstream

sectoral paths already internalise the same end-use efficiency savings.

5 Empirical validation of forecasts

This appendix documents the empirical validation exercises used to calibrate the industry-level productivity mapping and to quantify forecast uncertainty. The validation has three components. First, we test the fixed-weight Domar aggregation approximation used in the forecasts. Second, we calibrate and validate the industry-level mapping from output and real prices to industry-level TFP growth. Third, we describe how the empirical forecast errors are propagated into the scenario uncertainty bands.

Throughout the appendix and the main text (Fig. 1), forecast errors are defined as predicted minus actual. For example, if $\widehat{A}^{\text{pred}}$ denotes a forecast and $\widehat{A}$ the realised value, the reported error is

$$\varepsilon = \widehat{A}^{\text{pred}} - \widehat{A}. \quad (35)$$

When constructing prediction bands, this convention implies that the forecast error is added with a minus sign: a realised value can be written as $\widehat{A} = \widehat{A}^{\text{pred}} - \varepsilon$. Equivalently, the additive correction to a point forecast is $-\varepsilon = \widehat{A} - \widehat{A}^{\text{pred}}$.

Unless stated otherwise, the validation exercises pool observations across countries, sectors (for the ζ validation), and forecast horizons τ , where τ denotes the length of the growth interval from t to $t + \tau$. This pooled sample mechanically contains many more short-horizon observations than long-horizon observations: for example, one-year growth rates can be observed for many adjacent year pairs, whereas the maximum 24-year horizon is observed only once for each country-sector time series.

A Data construction

The validation uses three datasets: EU KLEMS, OECD STAN, and NBER CES manufacturing. Across datasets, the relevant forecasting variables are cumulative industry TFP growth, cumulative real gross-output growth, and cumulative real output-price growth between a forecast origin t and a horizon $t + \tau$.

Real output-price growth is measured relative to an aggregate GDP or value-added deflator:

$$\hat{p}_{i,c,t,\tau}^r = \left[\ln P_{i,c,t+\tau}^Y - \ln P_{i,c,t}^Y \right] - \left[\ln P_{c,t+\tau}^{\text{GDP}} - \ln P_{c,t}^{\text{GDP}} \right]. \quad (36)$$

Thus, $\hat{p}_{i,c,t,\tau}^r > 0$ means that the industry's output price rises faster than the country-level GDP deflator over the same horizon. This is the empirical counterpart of the real-price adjustment discussed in Appendix 2. We implement this adjustment with the country aggregate value-added deflator rather than a sector-specific factor-price index. While Appendix 2 shows that a sector-specific deflator is more appropriate, a country-level deflator is much less noisy.

EU KLEMS. The KLEMS zeta-validation notebook constructs the working sample from the Growth Accounts and National Accounts files. It uses

$$Y_C = \text{GO_CP}, \quad Y_Q = \text{GO_Q}, \quad P_Y = \text{GO_PI},$$

$$X_C = \text{II_CP}, \quad X_Q = \text{II_Q}, \quad L_C = \text{LAB}, \quad L_Q = \text{LAB_QI}, \quad K_Q = \text{CAP_QI}.$$

We form gross-output cost shares L_C/Y_C , X_C/Y_C , and $1 - L_C/Y_C - X_C/Y_C$, then computes annual gross-output TFP growth from gross-output, labour, intermediate-input, and capital-quantity growth. Long-horizon TFP growth is computed from cumulated TFP levels, not by independently estimating a new residual at each horizon. The country deflator is the total-economy value-added price index, VA.PI, for sector TOT. We exclude aggregate or market-only sectors such as TOT, TOT_IND, MARKET, and MARKETxAG to ensure a coverage without overlap. We exclude EU11, EU12 for the same reason.

We exclude countries whose aggregate productivity or price series display large structural breaks or implausible movements over the validation window, namely Slovakia (SK), Ireland (IE), and Greece (EL).

OECD STAN. STAN is the OECD industry accounting database. It covers only the last years, It is therefore treated as a short-horizon robustness check, not as the baseline calibration. We use the following STAN variables:

$$Y_C = \text{PROD}, \quad Y_Q = \text{PRDK}, \quad P_Y = \text{PRDP}, \\ X_C = \text{INTI}, \quad X_Q = \text{INTK}, \quad V_C = \text{VALU}, \quad V_Q = \text{VALK}, \quad L_C = \text{LABR}, \quad L_Q = \text{HRSN}, \quad K_Q = \text{CPNK}.$$

If hours worked are missing for an entire country-sector series, we replace HRSN with total employment, EMPN, for that series. The country deflator is the aggregate value-added price index, VALP, for sector D01T99. The STAN zeta cross-check uses the available short-horizon support after dropping the aggregate sector and trimming the 5th and 95th percentiles of TFP growth, output-price growth, output growth, and aggregate inflation.

NBER CES. For NBER CES manufacturing, we use the published five-input TFP index:

$$A_{i,t} = \text{tfp5}.$$

Current-price output is `vship`, the industry output price index is `pishop`, and real gross output is

$$Y_Q = \frac{\text{vship}}{\text{pishop}}.$$

The GDP deflator is the BEA aggregate price index with code A191RG. We remove all NAICS industries beginning with 334, because computer and electronics industries dominate the sample through extreme price and TFP movements, and also remove sectors 325312 and 325314 (fertilizers) for the same reason.

B Validation of fixed-weight Domar aggregation

The first validation exercise asks whether aggregate TFP growth over horizon τ can be approximated by holding Domar weights fixed at the forecast origin. This is the relevant forecasting exercise because future expenditure shares are not observed when the forecast is made.

For country c , industry j , forecast origin t , and horizon τ , let $\hat{A}_{j,c,t,\tau}$ be cumulative industry gross-output TFP growth and let $\hat{A}_{c,t,\tau}$ be cumulative aggregate TFP growth. Fixed-origin Domar weights are

$$\lambda_{j,c,t} = \frac{\text{GO_CP}_{j,c,t}}{\text{VA_CP}_{c,t}^{\text{TOT}}}.$$

The “first-order”, fixed-weight prediction is

$$\hat{A}_{c,t,\tau}^{\text{pred},1} = \sum_j \lambda_{j,c,t} \hat{A}_{j,c,t,\tau},$$

with aggregation error

$$\varepsilon_{c,t,\tau} = \hat{A}_{c,t,\tau}^{\text{pred},1} - \hat{A}_{c,t,\tau}.$$

As an ex-post robustness check, we also compute a second-order approximation using average initial and terminal Domar weights,

$$\hat{A}_{c,t,\tau}^{\text{pred},2} = \sum_j \frac{\lambda_{j,c,t} + \lambda_{j,c,t+\tau}}{2} \hat{A}_{j,c,t,\tau}.$$

This second-order version is not used in the forecast baseline because it requires terminal weights, but it is useful for checking whether fixed initial weights are the main source of aggregation error.

The final KLEMS Domar validation sample keeps countries with complete first-digit sector coverage. We thus exclude Spain, Lithuania, the United Kingdom, Luxembourg, and the United States, which leaves 18 countries




Figure 4: Distribution of fixed-weight Domar-aggregation errors by forecast horizon in EU KLEMS. The solid line is the pooled mean of $\varepsilon_{c,t,\tau}$, and the shaded bands report the empirical 25–75 percent and 10–90 percent ranges across countries and forecast origins.

and horizons up to 12 years, and 11 afterwards. We exclude horizons longer than 20 years as the number of observations drops (see Fig. 6). The pooled first-order aggregation error is close to zero at short and medium horizons, while dispersion increases with the horizon (Fig. 4). At horizon 20, the recomputed first-order mean error is 0.3pp, with a 10–90 percent range from -2.9 to 4.0 pp.

Figure 5 shows that allowing terminal weights to enter the ex post approximation does not materially reduce the error bands. We therefore retain the first-order fixed-weight approximation in the forecast and use its empirical error distribution for uncertainty calibration.

The diagnostic in Fig. 6 shows a marked loss of support after a 20-year horizon: the number of observations declines, and the country coverage also drops. We therefore estimate the empirical error distribution up to $\tau = 20$ and linearly extrapolate the error-band trends beyond that point, as shown in Fig. 7. This avoids letting few observations in the long horizons determine the long-horizon bands. We cannot rule out genuine long-horizon bias or data artefact, we thus choose to reduce sensitivity to a part of the sample where data coverage is weak.

C Validation of the industry-level mapping from prices and output to industry TFP growth

The main text uses the industry-level relationship

$$\hat{A}_{i,c,t,\tau}^{\text{pred}}(\zeta) = \zeta \hat{X}_{i,c,t,\tau} + (\zeta - 1) \hat{p}_{i,c,t,\tau}^r, \quad (37)$$

where $\hat{X}_{i,c,t,\tau}$ is real gross-output growth and $\hat{p}_{i,c,t,\tau}^r$ is real output-price growth as defined in Eq. 36. For any candidate ζ , the industry-country-origin-horizon-level forecast error is

$$v_{i,c,t,\tau}^{\zeta}(\zeta) = \hat{A}_{i,c,t,\tau}^{\text{pred}}(\zeta) - \hat{A}_{i,c,t,\tau}.$$

The baseline calibration chooses ζ to minimise mean absolute error,

$$\text{MAE}(\zeta) = \frac{1}{N} \sum_{i,c,t,\tau} |v_{i,c,t,\tau}^{\zeta}(\zeta)|.$$

Table 5 reports the values of ζ under several specifications. The baseline is KLEMS at horizons $\tau \geq 15$, which gives $\zeta = 0.128$. The all-horizon KLEMS value is higher because short-horizon observations are much




Figure 5: First- and second-order Domar-aggregation errors by forecast horizon in EU KLEMS. Blue denotes the first-order fixed-initial-weight approximation; red denotes the second-order approximation using the average of initial and terminal weights.

more numerous and reflect more short-run price-output comovement. NBER CES gives a very similar long-horizon value after removing the computer/electronics and fertilizer outliers. STAN is higher, but it has only short-horizon support and is noisier after trimming.

	Sample	Horizons	Observations	ζ_{MAE}
KLEMS	Baseline long-horizon sample	15–24	25,913	0.128
KLEMS	Reduced long-horizon sample	15–20	22,056	0.135
KLEMS	All horizons	1–24	178,482	0.171
KLEMS	Reduced long-horizon sample	1–20	174,625	0.175
STAN	5–95 percent trimmed sample	1–8	29,911	0.190
NBER CES	Clean long-horizon sample	15–30	195,685	0.133
NBER CES	Clean all-horizon sample	1–58	573,755	0.130

Table 5: ζ calibrations from the validation datasets. The objective is mean absolute industry-level TFP forecast error.

As in Fig. 6 for the Domar aggregation, Fig 8 shows that the empirical support for estimating the ζ forecast-error distribution declines steadily with the horizon, with a sharper loss of observations after about $\tau = 20$, when we have at least 2000 observations. This makes the long-horizon quantiles increasingly sensitive to sample composition and outlying country-sector-year cells. For this reason, we use the empirical distribution up to $\tau = 20$ and extrapolate the error-band trends beyond that point, rather than letting the sparsely observed terminal horizons drive the confidence intervals.

Figure 9 summarises the forecasting accuracy of the industry-level mapping in the three validation datasets. In all datasets, forecast dispersion increases with the horizon, which is why the scenario analysis uses empirical prediction intervals rather than a single point forecast.

Prediction intervals. We estimate the empirical distribution of errors on industry-level TFP growth at each horizon and report the 10th, 25th, 50th, 75th, and 90th percentiles.

The observed KLEMS support reaches 24 years; we extend the quantiles to horizon 30 by fitting a linear trend to the last 10 observed horizons for each quantile. The horizon-30 additive KLEMS correction has median error of -0.0384 and a 10–90 percent range from -0.2960 to 0.1501 in log.




Figure 6: Support for estimating Domar aggregation error distributions at long horizons. Beyond a 20-year horizon, both the number of country observations and the number of countries represented in the panel fall sharply. The apparent downward drift in the estimated error distribution at longer horizons may therefore reflect sparse-sample composition rather than a genuine horizon-dependent bias.

D Robustness to the calibration objective

The baseline value of ζ minimises mean absolute forecast error, but the calibration is not meant to depend on a single arbitrary loss function. Table 6 therefore recomputes ζ under several objectives. Accuracy-oriented objectives give similar values in the main validation datasets. In KLEMS, mean absolute error, median absolute error, root mean squared error, and output-share weighted mean absolute error imply ζ values between 0.119 and 0.150. In NBER CES, the corresponding values lie between 0.112 and 0.146. STAN is less stable, as expected given its shorter horizon support and noisier sample, but its accuracy-oriented objectives remain in the same broad range.

Let o index an observation in the pooled validation sample $\mathcal{O}$. For the last two rows of table 6, we define

$$\zeta^{\text{output}} = \arg \min_{\zeta \in [0,1]} \frac{\sum_{o \in \mathcal{O}} s_{i,o} |v_o(\zeta)|}{\sum_{o \in \mathcal{O}} s_{i,o}}, \quad \zeta^{\text{share}} = \arg \min_{\zeta \in [0,1]} \frac{\sum_{o \in \mathcal{O}} (s_{i,o}/s_o) |v_o(\zeta)|}{\sum_{o \in \mathcal{O}} (s_{i,o}/s_o)}.$$

Table 6: Values of ζ minimising forecasting error under several metrics and in three databases. The KLEMS columns report both the all-horizon sample and the long-horizon sample with horizon ≥ 15 and $15 \leq \text{horizon} \leq 20$; for STAN, outliers are excluded; for NBER CES, $15 \leq \text{horizon} \leq 30$ and outlier industries are excluded.

Error metric	KLEMS all horizons	KLEMS horizon ≥ 15	KLEMS $15 \leq \text{horizon} \leq 20$	STAN	NBER CES
<i>Bias-based metrics</i>					
Mean error	0.127	0.129	0.137	0.000	0.012
Absolute mean error	0.127	0.129	0.137	0.000	0.012
Median error	0.107	0.139	0.148	0.000	0.028
<i>Absolute-error metrics</i>					
Mean absolute error	0.171	0.128	0.135	0.190	0.133
Median absolute error	0.180	0.150	0.153	0.145	0.112
Mean absolute output weighted error	0.333	0.233	0.244	0.357	0.124
Mean absolute output-share weighted error	0.159	0.119	0.125	0.132	0.139

Continued on next page

Error metric	KLEMS all horizons	KLEMS horizon ≥ 15	STAN	NBER CES
<i>Squared-error metrics</i>				
Root mean square error	0.193	0.137	0.148	0.219 0.146

E Econometric validation of ζ in KLEMS

The calibration above treats ζ as the scalar that minimises out-of-sample-style backtest errors. Here we estimate ζ using standard regression techniques, which make it easier to test whether specific periods or countries drive the average result, by including fixed effects. We can do this by noticing that, from e.g. Eq. 37, we can rearrange the definition of ζ as

$$\hat{A}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r = \zeta (\hat{X}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r).$$

We estimate the regression

$$\hat{A}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r = \zeta (\hat{X}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r) + \xi_{i,c,t,\tau},$$

where $\xi_{i,c,t,\tau}$ is a noise term. Table 7 shows the results, where the baseline does not include an intercept, and subsequent columns include various combinations of fixed effects. Here we have chosen KLEMS at horizon 18, which is close to the long-horizon calibration window but still has 3,345 observations. The estimated coefficient is 0.133 without fixed effects and 0.129 with country, sector, and year fixed effects.

Table 7: Regression Models with Different Fixed Effects, KLEMS at horizon 18 years

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Dependent variable: $\hat{A}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r$</i>							
$\hat{X}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r$	0.133***	0.137***	0.136***	0.133***	0.136***	0.134***	0.131***
Std Error	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Country FE		✓			✓		✓
Sector FE			✓		✓		✓
Year FE				✓		✓	✓
R^2	0.138	0.185	0.525	0.142	0.568	0.190	0.573
N	3345	3345	3345	3345	3345	3345	3345

The fixed-effect estimates are very close to the baseline MAE calibration, indicating that the calibrated value is not driven by persistent country, sector, or year components that could mechanically generate a positive price-output-productivity relationship.

For Figure 11, we estimate the parameter separately for each forecast horizon τ . Concretely, for each τ we restrict the sample to observations with that horizon and estimate

$$\hat{A}_{i,t,t+\tau} + \hat{p}_{i,c,t,\tau}^r = \zeta_\tau (\hat{X}_{i,t,t+\tau} + \hat{p}_{i,c,t,\tau}^r) + \alpha_i + \gamma_c + \delta_t + \xi_{ict\tau},$$

so Panel A plots the OLS estimate ζ_τ for each horizon and database. Panel B is constructed analogously: for each horizon τ , we restrict the sample to that horizon and choose the value of ζ_τ that minimises the Mean Absolute Error (MAE) of the implied prediction for $\hat{A}_{i,t,t+\tau}$. Hence both panels report horizon-specific estimates.

Figure 11 shows that regression estimates are larger at short horizons, where cyclical price and quantity movements are more prominent, but settle close to 0.13 at medium-to-long horizons. Recomputing selected fixed-effect regressions gives coefficients of 0.195 at horizon 15, 0.129 at horizon 18, 0.132 at horizon 20, and 0.130 at horizon 23. This supports using $\zeta \simeq 0.13$ for the long-run scenario applications. Note that NBER CES is limited to Manufacturing sectors. STAN is limited to short time series. We nevertheless show that the three database give close estimates of the value of ζ .




Figure 7: Linear extrapolation of Domar aggregation error bands beyond the 20-year empirical cutoff.

F Combining industry-level and aggregation uncertainty

The aggregate forecast combines two empirical error distributions: the industry-level error from using our empirical industry-level mapping and ζ and the aggregation error from fixed-weight Domar aggregation. Let $\mathcal{I}$ be the set of sectors included in the analysis. With the appendix sign convention, the total aggregate forecast error is

$$\rho_{c,t,\tau} = \frac{\sum_{j \in \mathcal{I}} \lambda_{j,c,t}}{\sum_i \lambda_{i,c,t}} \varepsilon_{c,t,\tau} + \sum_{j \in \mathcal{I}} \lambda_{j,c,t} v_{j,c,t,\tau}.$$

Equivalently, since realised values satisfy

$$\widehat{A} = \widehat{A}^{\text{pred}} - \rho,$$

the prediction bands are constructed by adding the opposite-signed residual $-\rho$ to the point forecast. The same accounting identity is used in the scenario fan charts, with empirical quantiles sampled by horizon.

The simulations draw one Domar aggregation error and one industry-level TFP forecast residual per forecasted sector and simulation draw. The baseline treats sector-level ζ residuals as independent conditional on the horizon. This transparent specification is an approximation based on the independence assumption of TFP shocks and that errors are uncorrelated across industries.

G Separating Sectoral Calibration from Domar Aggregation

As shown in SI 2B, there exists, for any historical realisation of sectoral prices, quantities, and Domar weights, a value of ζ that exactly recovers aggregate productivity. Equivalently, combining Eq. 24 with the Domar aggregation formula in Eq. 25 gives the full forecasting equation

$$\widehat{A}_{i,\tau}^{\text{pred}}(\zeta) = \sum_i \lambda_{i,t} \left[\zeta \hat{X}_{i,t,\tau} + (\zeta - 1) \hat{p}_{i,t,\tau} \right],$$

where $\lambda_{i,t} = \text{Sales}_{i,t} / \text{GDP}_t$ is the Domar weight.

A natural way to validate the full pipeline would therefore be to choose ζ directly from the aggregate




Figure 8: Support for estimating the industry-level TFP forecast error distributions at long horizons. Beyond a 20-year horizon, both the number of country observations and the number of countries represented in the panel fall sharply. The apparent downward drift in the estimated error distribution at longer horizons may therefore reflect sparse-sample composition rather than a genuine horizon-dependent bias.

forecasting error,

$$\hat{A}_{t,\tau}^{\text{pred}}(\zeta) - \hat{A}_{t,\tau}.$$

This would amount to a Domar-weighted calibration of the industry-level mapping. We do not use this as the baseline, not because the full equation is invalid, but because this objective changes the estimand. It calibrates ζ to fit historical aggregate outcomes under historical Domar weights, whereas the parameter needed in the scenario analysis is the sector-level mapping from price and quantity paths into productivity growth.

The distinction matters for both point estimates and uncertainty bands. A full-pipeline calibration can fit aggregate data well by exploiting historical cancellations across sectors: positive sectoral errors in some industries can offset negative errors in others once multiplied by Domar weights. But the transition scenarios need not reproduce the historical pattern of sectoral shocks or those historical cancellations. Our uncertainty propagation instead treats industry-level errors separately from Domar aggregation errors, consistent with the maintained approximation that sector-level TFP shocks are independent conditional on horizon. A Domar-weighted calibration would blur these two objects and could make the empirical confidence intervals too narrow or centred on the wrong sectors.

For sector i , let the true industry-level coefficient be ζ_i and write

$$\hat{A}_i = \zeta_i \hat{X}_i + (\zeta_i - 1) \hat{p}_i.$$

Using a universal ζ , defined in Eq. 29, gives the forecast

$$\hat{\hat{A}}_i(\zeta) = \zeta \hat{X}_i + (\zeta - 1) \hat{p}_i,$$

so the sector-level error is

$$v(\zeta) = \hat{\hat{A}}_i(\zeta) - \hat{A}_i = (\zeta - \zeta_i)(\hat{X}_i + \hat{p}_i) = (\zeta - \zeta_i) \hat{s}_i, \quad (38)$$

with $\hat{s}_i \equiv \hat{X}_i + \hat{p}_i$ the combined sectoral price-and-quantity impulse. A Domar-weighted calibration therefore pulls ζ toward sectors with large historical $\lambda_i |\hat{s}_i|$, even if those are not the sectors driving the future transition scenario.




Figure 9: Forecast accuracy of the industry-level mapping by database. Each panel reports mean absolute error and root mean squared error of industry-level TFP predictions by forecast horizon.

Two-sector example. Suppose the economy has manufacturing M and services S , with $\lambda_S \gg \lambda_M$. Assume the net-zero transition affects manufacturing but not services, so $\hat{s}_S = 0$ in the application and $\hat{s}_M \neq 0$. The aggregate forecast error in that application is

$$\rho(\zeta) = \lambda_M(\zeta - \zeta_M)\hat{s}_M.$$

A full-economy weighted historical calibration would be pulled toward ζ_S because services dominate output. When the weighted and unweighted estimators differ, a weighted estimator can fit historical aggregate data well and yet perform worse for the transition scenario, simply because the shock is concentrated in a smaller set of sectors.

Baseline calibration choice. This is why our baseline calibration minimises sector-level forecast error across observations rather than weighting sectors by Domar weight: the baseline targets the industry-level mapping itself. Aggregate uncertainty is then handled separately through the Domar-validation exercise and the propagation of empirical errors. For completeness, we also report below what the full-pipeline calibration would imply if one nevertheless chose ζ directly from aggregate KLEMS backtesting errors.

Full-pipeline calibration: empirical values. Table 8 reports the resulting values, and Fig. 12 shows the corresponding aggregate error bands. Using observed initial Domar weights, the long-horizon aggregate mean-absolute-error criterion gives $\zeta = 0.101$, below the sector-level KLEMS baseline of $\zeta_{\text{industry}} = 0.128$. As a diagnostic, the lower panel of the table sets all Domar weights to one while still evaluating the aggregate forecast; this counterfactual gives larger values, $\zeta \approx 0.18$ under long-horizon MAE and RMSE. The next three paragraphs interpret this pattern.

Table 8: Values of ζ minimising aggregate KLEMS backtesting errors. The first panel uses observed initial Domar weights. The second panel is an aggregate counterfactual in which every Domar weight is set to one before summing sector predictions.

Sample	Metric	ζ
All horizons	Absolute mean error	0.096
All horizons	Mean absolute error	0.133
All horizons	Absolute median error	0.102

All horizons	Median absolute error	0.163
All horizons	Root mean square error	0.122
Long horizons ($h \geq 15$)	Absolute mean error	0.107
Long horizons ($h \geq 15$)	Mean absolute error	0.101
Long horizons ($h \geq 15$)	Absolute median error	0.099
Long horizons ($h \geq 15$)	Median absolute error	0.105
Long horizons ($h \geq 15$)	Root mean square error	0.099
<i>Aggregate counterfactual with all $\lambda_i = 1$</i>		
All horizons	Mean absolute error	0.184
All horizons	Root mean square error	0.190
Long horizons ($h \geq 15$)	Mean absolute error	0.180
Long horizons ($h \geq 15$)	Root mean square error	0.184

Decomposing the unit- λ counterfactual. The unit- λ rows separates two distinct distortions of the full-pipeline calibration. To see this, note that the aggregate forecast error under generic sector weights w_i has a simple form, with v_i defined in Eq. 38,

$$E(\zeta; w) = \sum_i w_i v(\zeta) = \underbrace{\zeta \sum_i w_i \hat{s}_i}_{\equiv S(w)} - \underbrace{\sum_i w_i \zeta_i \hat{s}_i}_{\equiv Z(w)},$$

and the aggregate calibration finds the ζ that makes $\zeta S(w) - Z(w)$ small on average. First, note E is a signed sum of sectoral errors: a positive v_i in one industry can be cancelled by a negative v_j in another before any squaring or absolute value is taken. The sector-level objective $\sum_i v_i(\zeta)^2$, by contrast, sums positive quantities and admits no such cancellation. Second, the weights w_i change which sectors carry the most influence: replacing $w_i = 1$ by $w_i = \lambda_i$ reshapes both $S(w)$ and $Z(w)$ and pulls the minimiser toward sectors with the largest Domar weight, hence historical sales.

Reading the two rows. The two effects can be turned on independently. Setting $w_i = 1$ kills the second distortion while leaving the first intact: cancellations across sectors are still active in E , but no sector is privileged by its sales share. The move from $\zeta_{\text{industry}} = 0.13$ to the unit- λ value of $\zeta \approx 0.18$ therefore isolates the *cancellation channel*. Switching back from $w_i = 1$ to observed $w_i = \lambda_i$ isolates the *Domar-composition channel*, and pulls the estimate down to $\zeta \approx 0.10$. Both shifts are around 0.05–0.08, small relative to the value of ζ , and they have opposite signs. The unit- λ diagnostic makes the resulting biases visible and quantifiable: roughly 0.08 in ζ units reflects how historical sales shares happen to weight sectoral errors, while about ~ 0.05 reflects cross-sector cancellations within the aggregation. It tells us the historical full-pipeline value is the sum of three things — the sector-level mapping (0.13) plus a cancellation contribution (0.05) plus a Domar-composition contribution (-0.08)

Why this strengthens the sector-level baseline. A scenario whose Domar footprint differs from the calibration sample—the natural case for a net-zero transition concentrated in a small set of high-emitting industries. The full-pipeline ζ is therefore tuned to a sample-specific cancellation that the scenario does not inherit, and the sample-specific Domar correction that partly offset that cancellation is also absent. The sector-level baseline avoids both by construction: it targets the mapping directly, and Domar aggregation is handled in a separate step that re-introduces λ and $\hat{s}_i$ at the application time, under the maintained independence assumption that governs the uncertainty propagation.




Figure 10: Empirical KLEMS prediction intervals for industry-level TFP forecasts, with $\zeta = 0.13$. The dark and light bands report the 25–75 percent and 10–90 percent ranges of the additive residual $\nu = \hat{A}^{\text{pred}} - \hat{A}$. The dashed vertical line marks the start of the extrapolated segment.

Part II

Applications

6 Climate Change Committee scenarios

Data and notation. We use the Climate Change Committee (CCC) Sixth Carbon Budget dataset (version 2) [11] for five scenarios and 14 sectoral categories over 2020–2050. For each sector i and year t , the dataset reports additional capital expenditure, additional operating expenditure, and annualised resource cost relative to the CCC baseline. We denote by $C_{i,t}$ the annual additional cost of the Net Zero transition, the CCC annualised resource cost in the baseline specification.

Implementation details. Before aggregating sectoral costs, we remove workbook summary rows such as “Total” and ensure that sectoral coverage is harmonised across sheets.

A Framework

The basic idea is simple. The net-zero transition changes production costs through additional capital expenditure, additional operating expenditure, and the operating savings that some investments generate. The CCC reports these objects for 14 sectors and five scenarios, and combines them into a time series of annualised resource costs, denoted $C_{i,t}$.² We interpret $C_{i,t}$ as the additional annual resource cost borne by sector i at date t , relative to a baseline without the transition.

We assume that sectoral output volumes remain constant. In CB6, the CCC explicitly states that the pathways “deliver the same services [as the baseline] without emissions” [29, p. 242]. and CB7 states “output quantities for all subsectors are the same as in the baseline” [30, p. 228]. We therefore treat the CCC costs, in both CB6

²See the CCC overview of the Sixth Carbon Budget numbers and the associated dataset.




Figure 11: Estimates of ζ by forecast horizon from regressions of $\hat{A}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r$ on $\hat{X}_{i,c,t,\tau} + \hat{p}_{i,c,t,\tau}^r$, with country, sector, and year fixed effects and optimisation of ζ to minimise the Mean Absolute Error (MAE) over the 3 databases for each horizon.

and CB7, as the cost of producing a comparable sectoral output path with lower emissions, while abstracting from volume changes, since they are not reported consistently across sectors. This assumption isolates the cost channel and provides a transparent benchmark; we assess its importance below.

Let $C_{i,t}$ denote the additional annual resource cost in sector i at date t . The change in this annual cost wedge between a base year t and horizon $t + \tau$ is

$$\Delta C_{i,t,\tau} = C_{i,t+\tau} - C_{i,t}.$$

Dividing by base-year output X_i converts this change into a change in unit cost. Assuming that changes in unit costs pass one-for-one into prices, the corresponding price change is approximated by

$$\hat{p}_{i,t,\tau} \equiv \frac{p_{i,t+\tau} - p_{i,t}}{p_{i,t}} \approx \frac{\Delta C_{i,t,\tau}/X_i}{p_{i,t}} = \frac{\Delta C_{i,t,\tau}}{p_{i,t}X_i}.$$

This approximation treats $\hat{p}_{i,t,\tau}$ as a proportional change rather than the log change used in the theoretical framework.³

It is important not to sum annual costs over time. The object $C_{i,t}$ is an annual flow cost at date t , not an increment to a cumulative price wedge. What matters for prices, and therefore for TFP, is the change in the annual cost wedge between two dates. For example, suppose a sector incurs an additional cost of \$1M per year from year $t+1$ onward, and that this cost then remains constant. Prices rise once relative to the baseline to reflect the higher unit cost, implying a lower level of TFP. But if the annual cost remains at \$1M thereafter, prices do not keep rising, and productivity does not keep falling. Summing annual costs across years would therefore incorrectly treat a constant cost wedge as if it generated ever-larger price distortions over time. The relevant

³A more exact expression would be

$$\hat{p}_{i,t,\tau} = \log\left(1 + \frac{\Delta C_{i,t,\tau}}{p_{i,t}X_i}\right).$$

The proportional approximation is accurate only for sufficiently small cost changes.




Figure 12: Aggregate KLEMS backtesting errors when the industry-level ζ mapping and fixed-origin Domar aggregation are validated jointly. The point forecast uses the long-horizon aggregate MAE calibration, $\zeta = 0.101$. The solid line is the mean error, the dashed line is the median error, and the shaded regions show the empirical 25–75 and 10–90 percent ranges by horizon. The vertical dashed line marks the start of the linear extrapolation beyond the empirical cutoff at $\tau = 20$.

object is the change in annual costs, $\Delta C_{i,t,\tau}$, not the sum of past annual costs.

Using Eq. 24 and imposing zero output growth, $\hat{X}_{i,t,\tau} = 0$, the implied change in sectoral TFP between t and $t + \tau$ is

$$\hat{A}_{i,t,\tau} = \zeta \hat{X}_{i,t,\tau} + (\zeta - 1) \hat{p}_{i,t,\tau} = (\zeta - 1) \hat{p}_{i,t,\tau} \approx (\zeta - 1) \frac{\Delta C_{i,t,\tau}}{p_{i,t} X_i}.$$

Aggregating with Domar weights (Eq. 15) gives

$$\begin{aligned} \hat{A}_{t,\tau} &= \sum_{i=1}^N \frac{p_{i,t} X_i}{p_t Y_t} \hat{A}_{i,t,\tau} \\ &\approx \sum_{i=1}^N \frac{p_{i,t} X_i}{p_t Y_t} (\zeta - 1) \frac{\Delta C_{i,t,\tau}}{p_{i,t} X_i} \\ &= (\zeta - 1) \frac{\sum_{i=1}^N \Delta C_{i,t,\tau}}{p_t Y_t}, \end{aligned} \tag{39}$$

where $p_t Y_t$ is nominal GDP in the base year. This expression gives the change in aggregate TFP between t and $t + \tau$: aggregate productivity growth depends on the change in annual transition-related costs over that interval, expressed as a share of base-year GDP.

Below, we compute aggregate productivity growth in CCC scenarios in two ways. First, we use the aggregate “Whole economy” resource cost reported by the CCC. Second, we decompose the result into sectoral contributions. Under the sectoral approach, sector i ’s contribution to aggregate productivity growth is

$$\lambda_i \hat{A}_i = (\zeta - 1) \frac{\Delta C_{i,t,\tau}}{p_t Y_t}.$$

B Results

Table 9 reports the 2022–2050 TFP effects of the CCC CB6 scenarios and the CB7 Balanced Pathway, expressed in percentage points. The CB7 Balanced Pathway implies a small positive aggregate TFP effect of 0.12 pp. By contrast, all five CB6 scenarios imply aggregate productivity losses, ranging from -0.6 pp in Widespread Engagement to -1.3 pp in Tailwinds; the CB6 Balanced Net Zero Pathway lies in between, at -0.6 pp.

The sectoral decomposition in Table 9 shows that the negative CB6 aggregate effects are mainly driven by residential buildings, removals, manufacturing, non-residential buildings, and shipping. These losses are only partly offset by surface transport, whose contribution is positive in all CB6 scenarios, between +0.35 pp and +0.52 pp. The CB7 Balanced Pathway changes the balance: surface transport contributes +1.15 pp, large enough to offset sizeable negative contributions from residential buildings (-0.57 pp), removals (-0.25 pp), and non-residential buildings (-0.19 pp), yielding a small positive aggregate effect of +0.12 pp. Electricity supply is not uniformly negative: it contributes +0.21 pp in the CB6 Balanced Net Zero Pathway and +0.13 pp in CB7, but remains negative in Headwinds (-0.07 pp), Tailwinds (-0.35 pp), and Widespread Innovation (-0.08 pp). This contrast suggests that the aggregate sign is highly sensitive to sector-level assumptions, especially for surface transport, buildings, removals, and electricity supply, motivating the construction of transition scenarios based directly on technology deployment and cost assumptions.

Table 9: 2022–2050 TFP impact of CCC CB6 and CB7 Net Zero transition pathways, in percentage points

Sector	CB7 – Balanced Pathway	CB6 – Balanced Net Zero Pathway	CB6 – Headwinds	CB6 – Tailwinds	CB6 – Widespread Engagement	CB6 – Widespread Innovation
Agriculture	0.00	0.00	0.00	0.00	0.00	0.00
Aviation	0.10	0.00	0.10	0.00	0.10	0.10
Electricity supply	0.10	0.20	-0.10	-0.40	0.00	-0.10
F-gases	0.00	0.00	0.00	-0.00	-0.00	0.00
Fuel supply	-0.10	0.10	0.10	0.10	0.20	0.00
LULUCF sinks	-0.00	-0.10	-0.10	-0.10	-0.10	-0.10
LULUCF sources		0.00	0.00	0.00	0.00	0.00
Manufacturing	0.00	-0.20	-0.10	-0.10	-0.20	-0.10
Non res buildings	-0.20	-0.10	-0.20	-0.10	-0.10	-0.10
Removals	-0.30	-0.20	-0.30	-0.50	-0.20	-0.20
Res buildings	-0.60	-0.50	-0.60	-0.60	-0.50	-0.60
Shipping	-0.10	-0.10	-0.10	-0.10	-0.20	-0.10
Surface transport	1.20	0.40	0.40	0.50	0.40	0.50
Waste	-0.00	-0.10	-0.10	-0.10	-0.00	-0.10
Total	0.10pp	-0.60pp	-1.10pp	-1.30pp	-0.60pp	-0.80pp

C Accounting conventions

Bottom-up techno-economic scenarios are by nature constructed very differently from national accounts objects. There are different costs concepts available in the CCC scenarios, and here we use “annualised, processed resource costs”. While it is not possible to match perfectly these costs concepts to our national accounting framework, we briefly explain our choices.

Allocation of system costs to industries. “Processed” costs means costs after the CCC’s internal cost reallocation. Before reallocation, the same underlying system cost may appear twice, e.g. once as CAPEX in electricity supply, and again as OPEX in the downstream sector through its electricity bill. This treatment is appropriate for aggregate accounting, because it prevents overlapping cost categories from being counted twice, but it creates a bias in industry-level costs because one needs to take a stance on which industry the cost should be attributed to.

Cost of capital and annualization. In the Sixth Carbon Budget, cost-of-capital assumptions were used in annualising capital expenditure, partly as a proxy for project risk. In the Seventh Carbon Budget, by contrast, private financing costs are excluded: capital expenditure is annualised using the Green Book 3.5% social discount rate [21]. For simplicity, we keep the costs as supplied by each version of the carbon budget, noting that the treatment of the cost of capital can make a difference, particularly for long-lived assets.

D Robustness

This exercise sets sectoral quantity growth to zero when applying the reduced-form mapping

$$\hat{A}_i = \zeta \hat{X}_i + (\zeta - 1) \hat{p}_i.$$

If modelled sectors expand in real terms, this may understate predicted productivity growth.

As a rough sensitivity check, suppose all modelled sectors experience cumulative real quantity growth equal to the OBR forecast for UK GDP per capita between 2022 and 2050 (33.8%). This is a crude assumption, as GDP per capita growth need not match sectoral gross-output growth. The omitted quantity term is then $\zeta \ln(1.338)$.

Using $\zeta = 0.13$, this implies about $0.13 \times \ln(1.338) \approx 0.038$, or 3.9 percentage points cumulatively. Setting $\hat{X}_i = 0$ therefore biases sectoral TFP downward in this sensitivity check.

At the aggregate level, this effect must be scaled by Domar weights. With a combined weight of about 0.5 (sectors A–F), this implies an upward shift of $0.5 \times 0.038 \approx 0.018$, i.e. 1.8 percentage points cumulatively. This is larger than the baseline CB7 estimate of 0.12%, and would be sufficient to make some scenarios slightly positive, although uncertainty bands would still include zero.

This calculation should be interpreted as a back-of-the-envelope bound rather than a realistic calibration, but it shows that the zero-quantity-growth assumption is not innocuous.

E Propagation of validation errors in the CCC exercise

Let $\hat{A}_{t,\tau}^{CCC}(\zeta)$ denote the aggregate CCC point estimate obtained from Eq. 39, evaluated at the chosen value of ζ . Let $\mathcal{I}$ be the set of CCC sectors. The uncertainty bands do not modify the CCC accounting itself; they add empirical validation errors to the point estimate.

For each horizon τ , we draw $v_{i,\tau}^{(b)}$ from the empirical horizon-specific distribution of the industry-level mapping error $v_{i,c,t,\tau}(\zeta) = [\zeta \hat{X}_{i,c,t,\tau} + (\zeta - 1) \hat{p}_{i,c,t,\tau}] - \hat{A}_{i,c,t,\tau}$ (Fig. 10) and $\varepsilon_{\tau}^{(b)}$ from the empirical distribution of the fixed-weight Domar-aggregation error $\varepsilon_{c,t,\tau} = \sum_j \lambda_j \hat{A}_{j,c,t,\tau} - \hat{A}_{c,t,\tau}$. As throughout the validation appendix, the sign convention remains the one used above: ε = predicted – realised. Hence the empirical correction added to a forecast is $-\varepsilon$ and $-\nu$.

Because ν is an industry-level TFP error, it must be converted into aggregate units using the fixed-origin Domar weight $\lambda_{i,UK,t}$ from Eq. B. In other words, even massive errors for very small sectors – with small Domar weights – should not cause big errors in the aggregate result. Since the CCC exercise covers only a subset of the UK economy, we also rescale the whole-economy Domar error by the coverage ratio,

$$\sum_{i \in \mathcal{I}} \lambda_i \simeq 0.31 \quad \text{and} \quad \frac{\sum_{i \in \mathcal{I}} \lambda_i}{\sum_j \lambda_j} \simeq 0.16,$$

the sectors covered by CCC roughly includes the sectors A to F and H, thus covering agriculture, mining, manufacturing, energy, transport, waste, and buildings.

For simulation draw b , the aggregate CCC path for the country c , here the UK, is

$$\hat{A}_{c,t,\tau}^{(b)} = \hat{A}_{c,t,\tau}^{CCC}(\zeta) - \sum_{i \in \mathcal{I}} \lambda_i v_{i,c,t,\tau}^{(b)} - \frac{\sum_{i \in \mathcal{I}} \lambda_i}{\sum_j \lambda_j} \varepsilon_{c,t,\tau}^{(b)}.$$

The baseline draws sector-level ζ errors independently within each horizon. This keeps the validation bands




Figure 13: CCC CB7 and CB6 aggregate TFP path with empirical validation bands. The solid line is the CB7 Balanced Pathway point estimate. The dark and light shaded areas report the 25–75 percent and 10–90 percent empirical bands obtained by propagating industry-level TFP and Domar-aggregation errors as described above.

transparent, but it abstracts from cross-sector correlation in forecast errors.

Numerical implication for the CCC bands. Applying this procedure to the CB7 Balanced Pathway gives a point estimate of 0.12% aggregate TFP growth in 2050. The empirical validation bands are wide relative to the point estimate: the 25–75 percent interval is $[-1.29\%, 0.45\%]$, and the 10–90 percent interval is $[-2.09\%, 1.17\%]$ (Fig. 13). The point estimate is therefore small relative to the empirical uncertainty implied by the historical validation exercises.

7 Energy

This appendix documents the energy-sector exercise used in the paper. We focus on the useful-energy results reported in the paper, together with the calibration and robustness checks that matter for interpretation.

Key concepts First, there is a distinction between final and useful energy: final energy is the processed form delivered to consumers and firms, while useful energy is what remains after accounting for conversion losses, representing the actual service (heat, kinetic energy, light, etc.). Second, the electricity sector must ensure market clearing at all times, ensuring supply meets demand. Additional solutions must be implemented in high renewables systems due to the inherent intermittence—variability in supply—of solar and wind. Storage captures excess energy for later use, and flexibility balances supply and demand through mechanisms like demand response and grid interconnections. Finally, the price of energy in the market is set by the most expensive production unit needed to meet demand at a given time, known as the marginal price. It results from a bidding process and reflects both capital expenditure (CAPEX) for infrastructure and operational expenditure (OPEX) for ongoing costs like fuel and maintenance. In our exercise, we assume that, on average, the electricity market clearing price is close to the average cost of electricity.

Abbreviations and terminology.

BEIS UK Department for Business, Energy & Industrial Strategy (functions transferred to DESNZ in 2023; legacy data series retain the BEIS label).

CCC Climate Change Committee, the UK’s statutory independent advisory body on climate policy.

CB6 / BNZP CCC Sixth Carbon Budget; Balanced Net Zero Pathway, the CCC’s central CB6 scenario.

Way et al. The global scenario workbook of Way, Ives, Mealy & Farmer [12], hereafter Way et al.

LCOE Levelised cost of energy: the constant per-MWh price that equates the discounted revenue of a generating asset to its lifetime CAPEX plus OPEX.

CAPEX / OPEX Capital expenditure (one-off investment cost) and operating expenditure (ongoing fuel, maintenance, and balancing costs).

CCGT Combined-cycle gas turbine.

CHP Combined heat and power (cogeneration).

CCS Carbon capture and storage.

P2X Power-to-X: conversion of electricity into storable energy carriers, principally hydrogen and derived synthetic fuels.

VRF Vanadium redox flow battery, a medium-duration storage technology suitable for multi-day cycling.

Mtoe / TWh / EJ Million tonnes of oil equivalent; terawatt-hour; exajoule. Conversions: 1 Mtoe $\approx$ 11.63 TWh; 1 EJ = 277.78 TWh.

We first describe our assumptions on prices, then our assumptions on quantities. Next, we present the results, followed by robustness checks.

A Input prices and cost assumptions

Long-run equivalence of price and average cost. Our productivity measures are built from technology cost paths rather than observed market prices. This relies on the assumption that, on average and over the horizon considered, the wholesale electricity price equals the long-run average cost of supply (LCOE, inclusive of CAPEX, OPEX, and a normal return on capital). In short-run electricity markets, clearing prices are set by the marginal unit and can diverge substantially from average cost, with the gap closed by scarcity rents, capacity

payments, and bilateral contracts that fund capital recovery. For a multi-decadal growth-accounting exercise, what matters is not within-year dispatch pricing but the resource cost of delivering energy services: under free entry and competitive capital markets, the time-averaged wholesale price must cover long-run average cost, and our TFP measures then track real resource savings rather than transfers between producers and consumers. This is the standard framing in the integrated-assessment and technology-forecasting literatures [12].

The assumption abstracts from two features of high-renewables systems that we address in reduced form elsewhere in the appendix: value deflation or “cannibalisation” of wind and solar revenues at high penetration [31], is partly captured by the 1.5× overgeneration markup on renewable prices; and network and balancing costs are captured by the grid-cost adder c_{grid} in the same section.

Data sources To ensure a fair comparison between scenarios, we use the same energy prices in all our scenarios (see Table 11 for a robustness check of this assumption).

All prices are expressed in constant 2020 USD/MWh, sources are summarised in Table 14. The price database combines forecast series from Way et al. [12] for oil, gas, coal, solar PV, wind, battery storage, and hydrogen with BEIS-based values for legacy electricity technologies (unabated gas electricity, gas+CCS electricity, biomass electricity, biomass+CCS electricity, nuclear, hydro, CHP gas, coal electricity, and bioenergy). Way et al. [12] forecast prices using AR(1) models for technologies that do not tend to make progress, and using Wright’s law, as defined in Lafond et al. [32], for technologies that do feature cost declines. Their forecasts imply substantial cost declines – as experience accumulates – for solar, wind, hydrogen and batteries but no trend for nuclear and fossil fuel based technologies.

For hydrogen, the price series follows the electrolyser cost path converted into USD/MWh of hydrogen. In the current code, hydrogen-based electricity is priced as hydrogen divided by a conversion efficiency of 0.5 plus a fixed 14 USD/MWh plant-cost adder based on the BEIS CCGT benchmark. For non-bio waste, the price path is interpolated linearly between the initial and terminal BEIS values.

For the paper results, wind and solar are treated as delivered electricity. Overgeneration is therefore represented as a higher effective cost per delivered MWh rather than as extra output. We implement this as a multiplicative markup of 1.5 on wind and solar prices, following Way et al. [12] calculations of total energy system costs (See Table 13 for robustness checks with a different overgeneration factor). For long-lived technologies, the baseline results use vintage prices: each year’s price is a weighted average of surviving build cohorts, with the vintage recursion anchored at 2022 so that 2022 vintage and marginal prices coincide by construction.

The transition of the energy systems requires investment in the electricity distribution network. First, because massive electrification means a scale up of the quantity of electricity that needs to be transmitted and thus the number of power lines, transformers, etc. Second, because the use of renewable electricity generation means that the decentralised production of electricity must be matched with the demand throughout the country. The grid must be able to accommodate such flexibility. Ref. [12] models the two components independently for the global energy system. We take a simpler approach and consider a renewable grid-cost adder of 7 USD/MWh [33] to wind and solar generation above the 2022 network baseline:

$$p_{e,t}^{\text{grid}} = p_{e,t} + c_{\text{grid}} \max\left(0, \frac{q_{R,t} - q_{R,2022}}{q_{R,t}}\right),$$

where $q_{R,t}$ is total wind-plus-solar output and $c_{\text{grid}} = 7$ in the baseline. The impact of the network cost on the UK Fast transition scenario is about 11pp on energy TFP by 2050.

A back of the envelope calculation gives a similar order of magnitude. Wind is about 20 USD/MWh in 2050, 32 USD/MWh when taking overgeneration into account. For solar, the cost is 10.6 USD/MWh, and 17 USD/MWh with overgeneration. A 7 USD 2020/MWh is an increase of about 21% for wind (44% of useful energy in 2050) and 40% for Solar (9% in 2050). That is an increase of about 12.5% of average energy price, with $\zeta = 0.13$, the global impact would be $(0.13 - 1) \times 0.125 = -0.109$, that is -11% on energy productivity.

We also report a robustness check with $c_{\text{grid}} = 14$ and $c_{\text{grid}} = 0$ (Table 11), and with all cost data from Way

et al. [12] for both the UK and the World scenario.

Table 10: Conversion factors between primary, final, and useful energy

Energy	Primary to final electricity	Final to useful
Gas	0.5	0.6
Oil	0.4	0.25
Bioenergy	0.4	0.25
Non-bio waste	0.4	0.25
Hydrogen	0.5	0.6
Solid/coal	0.4	0.6
Electricity	1	0.9

B Scenario Quantities

We use two quantity inputs. First, for the UK we use the Climate Change Committee Sixth Carbon Budget scenarios. Second, for the world we use the global scenario workbook from Way et al. [12] (Way et al.). UK quantities are expressed in TWh and world quantities in EJ.

For the CCC scenarios, the code harmonises the technology labels with the Way et al. taxonomy. Interconnection is excluded because it does not map cleanly to a production technology and is negative in some years. The CCC workbook reports both primary fuel use and final electricity generation; to avoid double counting, we net out fuel inputs used to generate electricity:

$$X_{\text{final fuel}} = X_{\text{primary fuel}} - \frac{X_{\text{electricity from fuel}}}{\beta},$$

where β is the primary-to-final conversion factor from Table 10.

We check the recovered energy production against the UK primary energy equivalent consumption in 2022 (170Mtoe Department for Business, Energy & Industrial Strategy [34]). The sum of the corrected quantities of energy in the Balanced Net Zero Pathway is 1866 TWh in 2022, which is about 160Mtoe (errors in conversion factors, density of energy, convention about curtailment and storage easily explain the 6% gap). Note that a 6% difference in energy quantity, would only result in a 0.8pp impact for the energy-industry TFP by 2050 ($\zeta \cdot \ln(0.06 + 1) = 0.13 \times 0.058 = 0.008$).

Electricity storage Electrifying the usages and increasing the share of renewable and intermittent energy (wind and solar) means increased needs for flexibility. Long-term storage is provided by P2X, whereas short-term daily storage is provided by Li-ion batteries and medium term multi-day storage is provided by flow batteries, namely Vanadium redox flow (VRF) (see appendix in [12], sections 3.1 and 3.2, for a review of the literature and the assumptions).

For the Way et al. scenarios, we add storage and hydrogen-electricity quantities explicitly following Way et al. [12]. Battery storage capacity is calibrated as 30% of generated wind-plus-solar electricity, where generated renewable electricity is taken to be 1.5 times delivered renewable electricity. This capacity is converted into annual energy by multiplying by 365, which preserves the annual cost accounting used in the appendix. Hydrogen-based electricity is taken from the Way et al. fast-transition and introduced as an explicit electricity source.

The CCC scenarios are not high-renewable, high-electrification scenarios (renewables are 25–49% of final energy by 2050), so storage plays only a limited role. Scenarios express storage needs directly in TWh (0.012 TWh in 2050 for the CB6 Balanced Net Zero Pathway).

The UK fast-transition scenario The scenario is built in three steps. First, the 2050 useful-energy total from the CCC Balanced Net Zero pathway is retained for the UK. Second, Way et al. fast-transition shares are applied to that UK 2050 total. Third, the 2050 solar share is reduced and reallocated to wind so that the UK path

matches the CCC 2050 wind-to-solar ratio of about 5; 76% of the Way et al. solar quantity is shifted to wind. Note that this is a conservative assumption as wind will be significantly more expensive than solar by 2050. The transition path between 2025 and 2050 is then interpolated linearly. The Way et al. scenario also relies heavily on hydropower, up to 5.5% of final electricity, more than what is available in the UK. We acknowledge the limitation and keep it as it is.

C Main useful-energy results

The paper focuses on useful energy. Figure 4 in the main article compares the three core scenarios used in the paper: the UK Balanced Net Zero pathway, a UK fast-transition counterfactual built from Way et al. fast-transition shares, and the global Way et al. Fast Transition.

Table 11: Useful-energy TFP gains and aggregate TFP gains, 2022–2050

Scen.	Price calib.	Grid (USD/MWh)	Domar wt.	ΔTFP (Energy)	ΔTFP (Agg.)
World Fast Transition	BEIS + Way et al.	0	0.063	165%	6.4%
World Fast Transition	BEIS + Way et al.	7	0.063	143%	5.8%
World Fast Transition	BEIS + Way et al.	14	0.063	124%	5.2%
World Fast Transition	All Way et al.	7	0.049	111%	3.7%
World Slow Transition	BEIS + Way et al.	7	0.063	69%	3.4%
World No Transition	BEIS + Way et al.	7	0.063	35%	1.9%
UK Fast Transition	BEIS + Way et al.	7	0.025	61%	1.2%
UK Fixed 2022 Energy Mix	BEIS + Way et al.	7	0.025	8%	0.2%
UK Balanced Net Zero Scenario	All Way et al.	7	0.022	9%	0.2%
CCC CB6 Balanced Net Zero Pathway	BEIS + Way et al.	7	0.025	13%	0.3%
CCC CB6 Headwinds	BEIS + Way et al.	7	0.025	17%	0.4%
CCC CB6 Widespread Engagement	BEIS + Way et al.	7	0.025	10%	0.2%
CCC CB6 Widespread Innovation	BEIS + Way et al.	7	0.025	21%	0.5%
CCC CB6 Tailwinds	BEIS + Way et al.	7	0.024	9%	0.2%

Range across UK CCC scenarios. The five CCC Sixth Carbon Budget scenarios, Fig. 14, illustrate the range of UK decarbonisation pathways considered in the CCC modelling. In our calibration, the implied energy-sector TFP gains range from 9% in Tailwinds to 21% in Widespread Innovation, with corresponding aggregate TFP gains between 0.2% and 0.5%. The Balanced Net Zero Pathway delivers a 13% energy-sector TFP gain and a 0.3% aggregate TFP gain under the BEIS + Way et al. price calibration.

D Counterfactual baselines

The productivity number we present in this section are the energy TFP growth from 2022 to 2050 under various energy transition scenarios. One could argue that the price of solar and wind will decrease no matter what the




Figure 14: Useful-energy comparison in the five CCC 6th Carbon Budget scenarios.

UK energy pathway looks like. We provide a quick back-of-envelop computation to compute what would TFP growth look under a fixed energy mix.

For the UK, let us assume the energy mix is fixed at its 2022 shares. We thus only take into account the change in total energy demand and the change in technology costs, without reallocation between energy. UK energy-sector TFP would rise by 8% over 2022–2050, implying an aggregate TFP gain of 0.2%.

For the world, we cannot consider that unit cost will decrease if the energy mix remains the same. The lower benchmark for the world energy TFP is then just due to the growth in volume ($\zeta\hat{X}$). Using the demand growth rate of approximately 2% per year reported in Way et al. [12], this counterfactual implies an increase of roughly 8% in global energy-sector TFP over 2022–2050, and about 0.5% in aggregate TFP gain. This provides a useful

lower benchmark for interpreting the Way et al. scenarios.

The No Transition scenario in Way et al. [12] should therefore not be interpreted as a no-productivity-growth baseline. Even without a fast transition, growing energy demand is partly met by technologies whose unit costs decline over time. In our calibration, the No Transition scenario generates a 35% increase in global energy-sector TFP and a 1.9% aggregate TFP gain. By contrast, the Fast Transition scenario under the BEIS + Way et al. calibration with a 7 USD/MWh grid cost generates a 143% increase in energy-sector TFP and a 5.8% aggregate TFP gain.

E Robustness to calibration assumptions

Calibration to Way et al. (2022) As a calibration check, we recompute global Fast Transition costs using Way et al. prices for all technologies and the same 1.5 overgeneration factor for wind and solar. Under this specification, total annual energy-system costs are USD 4.15 trillion in 2022 and USD 4.90 trillion in 2050, matching the order of magnitude reported by Way et al. [12]. When not considering the overgeneration of wind and solar total annual energy-system costs are USD 4.08 trillion in 2022 and USD 4.05 trillion in 2050, thus overgeneration represents 20% of 2050 total system costs. See Table 13 for the impact on TFP growth.

Alternative price calibration. Our baseline price calibration combines Way et al. prices for technologies with learning curves and BEIS assumptions for several legacy electricity technologies; see Table 14. Table 12 reports the unit costs that differ between this BEIS + Way et al. calibration and the alternative calibration using Way et al. prices for all available technologies. The main differences concern biomass electricity, nuclear, coal electricity, and gas electricity. Biomass electricity, coal electricity, gas electricity, and nuclear are substantially more expensive under the BEIS + Way et al. calibration.

Table 12: Unit-cost comparison for technologies whose BEIS and Way calibrations differ

Energy	BEIS + Way (USD/MWh)	Way (USD/MWh)
Nuclear	138.2	90.0
Hydro-electricity	102.8	103.0
Coal electricity	165.8	59.4
Biomass electricity	230.2	72.0
Gas electricity	69.3	54.9

Figure 15 repeats the useful-energy comparison after replacing the remaining legacy electricity prices with Way et al.-consistent values. This check addresses the concern that the main results might be driven by the hybrid calibration rather than by the scenario quantities and the relative price decline of renewables.

Using the all-Way et al. price calibration has modest effects on the UK Balanced Net Zero Scenario: the energy-sector TFP gain is 9%, compared with 13% for the CCC Balanced Net Zero Pathway under the BEIS + Way et al. calibration, and the aggregate TFP gain remains approximately 0.2–0.3%. The calibration choice matters more for the world fast-transition scenario. With BEIS + Way et al. prices and a 7 USD/MWh grid cost, the world fast transition implies a 143% gain in energy-sector TFP and a 5.8% gain in aggregate TFP over 2022–2050. Under all-Way et al. prices, the corresponding gains fall to 111% and 3.7%.

Higher grid-cost robustness. Table 11 repeats the main useful-energy comparison after removing and doubling the renewable grid-cost adder from 7 to 0 and 14 USD/MWh. This is a direct robustness check on the infrastructure-cost assumption embedded in wind and solar prices.

Overgeneration factor Table 13 repeats the main useful-energy comparison after removing the overgeneration factor of wind and solar. This factor represents that in order to produce 1 TWh of consumed electricity, the system actually produces 1.5 TWh on average, when the system is in overproduction, energy that is then stored in either batteries or hydrogen. Results are not qualitatively different.

Table 13: Sensitivity of useful-energy TFP gains to wind and solar overgeneration

Overgeneration assumption	Multiplier	Scenario	Domar wt.	Δ TFP (Energy)	Δ TFP (Agg.)
No overproduction	1.00	World Fast Transition	0.063	178%	6.6%
No overproduction	1.00	UK Fast Transition	0.024	88%	1.6%
No overproduction	1.00	UK Balanced Net Zero Scenario	0.024	24%	0.5%
Lower overproduction	1.25	World Fast Transition	0.063	159%	6.2%
Lower overproduction	1.25	UK Fast Transition	0.025	73%	1.4%
Lower overproduction	1.25	UK Balanced Net Zero Scenario	0.025	18%	0.4%
Current baseline	1.50	World Fast Transition	0.063	143%	5.8%
Current baseline	1.50	UK Fast Transition	0.025	61%	1.2%
Current baseline	1.50	UK Balanced Net Zero Scenario	0.025	13%	0.3%




Figure 15: Useful-energy comparison with Way et al.-consistent prices for the remaining legacy electricity technologies. **Note:** This figure uses the same scenario quantities as Figure 4 but replaces the remaining legacy electricity prices in the hybrid calibration with Way et al.-consistent values.

Table 14: Prices and sources per technology in the Energy sector

Technology (CCC)	Report	Ref	Techno	2020£/ MWh	USD 2020 /MWh	Horizon	Comments
Unabated gas electricity	BEIS 2023	table 10, p24 (in 2021 prices, p8 of BEIS 2023)	CCGT H Class	54	69.3	in 2025	Constant over time
H2GT electricity	BEIS 2023	table 10, p24 (in 2021 prices, p8 of BEIS 2023)	CCGT H Class (with- out H2 cost)	11	14.1	in 2025	Constant over time
Gas CCS electricity	BEIS 2020	table 4.3, p26	CCGT + CCS Post Combustion (FOAK)	82	112.3	in 2025	Constant over time
Biomass CCS electricity	BEIS 2020	table 8, p49, appendix 1	Biomass CCS (FOAK)	205	280.9	in 2025	Constant over time
Biomass electricity	BEIS 2020	table 8, p48, appendix 1	Biomass CHP	168	230.2	in 2025	Constant over time
Nuclear electricity	BEIS 2016 (take over in BEIS 2020)	table 4, p27	Nuclear PWR - FOAK	95	138.2	in 2025	Constant over time
Thermal low-carbon (Includes biomass and geothermal)	BEIS 2020	table 8, p48, appendix 1	Dedicated Biomass	98	134.3	in 2025	Constant over time
Other electricity (Includes coal, gas CHP, recip engines, other.)	BEIS 2020	table 8, p48, appendix 1	CCGT CHP mode	75	102.8	in 2025	Constant over time
Offshore Wind	Way et al. (2022)				39.60	in 2025	Learning curve
Onshore Wind	Way et al. (2022)				39.60	in 2025	Learning curve
Solar	Way et al. (2022)				52.21	in 2025	Learning curve
Offshore Wind	BEIS 2023	table 10, p24 (in 2021 prices, p8 of BEIS 2023)	Offshore wind	44	56.5	in 2025	(not used, for comparison)
Onshore Wind	BEIS 2023	table 10, p24 (in 2021 prices, p8 of BEIS 2023)	Onshore wind	38	48.8	in 2025	(not used, for comparison)
Solar	BEIS 2023	table 10, p24 (in 2021 prices, p8 of BEIS 2023)	Large solar	41	52.6	in 2025	(not used, for comparison)

Table 14 continued from previous page

Technology (CCC)	Report	Ref	Techno	2020£/ MWh	USD 2020 /MWh	Horizon	Comments
Other renewables (Includes hydro, tidal, and wave.)	BEIS 2020	table 8, p49, appendix 1	Hydro Large Storage	75	102.8	in 2025	Constant over time
Coal electricity (not in CCC scenario)	BEIS 2016	table 4, p27	Coal - ASC with oxy comb. CCS - FOAK (subtracted the CCS and carbon cost)	114	165.8	in 2025	Constant over time
Used curtailment (Reflects curtailed generation used for production of hydrogen. This is additional to other sources of generation.)	not included						
Interconnection	not included						
Storage (Includes batteries and pumped hydro, and reflects gross energy deployed into storage.)	Way et al. (2022)						Learning curve
Gas	Way et al. (2022)						No time trend
Petroleum/oil	Way et al. (2022)						No time trend
Solid fuel/coal	Way et al. (2022)						No time trend
Final non-bio waste	CCC CB6 2020	Non bio waste fuel, in aviation additional data, row 96, price of the Non-bio waste fuel in p/kWh HHV	Non bio fuel waste	62	84.9	2020 value	Linear decrease to 2050
Final non-bio waste	CCC CB6 2050	Non bio waste fuel, in aviation additional data, row 96, price of the Non-bio waste fuel in p/kWh HHV	Non bio fuel waste	47	64.4	2050 value	Linear decrease from 2020
Low carbon hydrogen	Hydrogen production cost BEIS (2021) Way et al. (2022)	BEIS 2020 hydrogen production cost - Annex, sheet 2020.E:	Alkaline electrolysis, Central case, Curtailed Electricity 25%, total LCOH	65	84	2020 value	Applying the learning curve of electrolysis from Way et al. (2022) to the cost of hydrogen (-76% from 2020 to 2050)

Table 14 continued from previous page

Technology (CCC)	Report	Ref	Techno	2020£/ MWh	USD 2020 /MWh	Horizon	Comments
Bioenergy/biomass	CCC CB6	Fuel Supply, M6.1, p20	table Domestic Biomass	22	30.1	2020 value	Constant over time